

Multi-Agent Non-Discriminatory Contracts*

Ke Ding¹, Bo Li¹, Ankang Sun²

¹Department of Computing, The Hong Kong Polytechnic University

²School of Computer Science and Technology, Shandong University

coco-ke.ding@connect.polyu.hk, comp-bo.li@polyu.edu.hk, ankang.sun@sdu.edu.cn

Abstract

We study multi-agent contracts, in which a principal delegates a task to multiple agents and incentivizes them to exert effort. Prior research has mostly focused on maximizing the principal’s utility, often resulting in highly disparate payments among agents. Such disparities among agents may be undesirable in practice, for example, in standardized public contracting or worker cooperatives where fairness concerns are essential. Motivated by these considerations, our objective is to quantify the tradeoff between maximizing the principal’s utility and equalizing payments among agents, which we call the *price of non-discrimination*. Our first result is an almost tight bound on the price of non-discrimination, which scales logarithmically with the number of agents. This bound can be improved to a constant by allowing some relaxation of the non-discrimination requirement. We then provide a comprehensive characterization of the tradeoff between the level of non-discrimination and the loss in the optimal utility.

1 Introduction

As one of the pillars in microeconomic theory, contract theory utilizes the principal-agent model to achieve optimal outcomes amidst asymmetric information [Holmström, 1979; Holmström, 1982]. It focuses on designing the market for services, analogous to how mechanism design and auction theory structure the market for goods. The past few years have witnessed growing interest in contract theory in both the economics and computation, with the applications extending to real-world scenarios like crowdsourcing platforms [Ho *et al.*, 2014], social media [Department, 2021], and government-run programs [Dai Li *et al.*, 2021].

In the vanilla hidden-action principal-agent model [Holmström, 1979; Grossman and Hart, 1992], the principal delegates a single task to an agent who may take one of the costly actions and its outcome determines the reward. A contract, structured as a monetary transfer, incentivizes the agent to

take the action maximizing his utility, which equals the transfer minus the cost. The principal cannot observe the action taken by the agent but only the outcome. Her goal is also to maximize her utility, which equals the reward obtained minus the payment transferred.

The basic model was later extended to the setting where the principal incentivizes a team of agents for a single project [Holmström, 1982; Babaioff *et al.*, 2006]. In this framework, each of the agents faces a binary choice: to exert effort at a specific cost, or shirk without cost. This multi-agent model involves strategic interactions among agents, ensuring that the resulting game induces a pure Nash equilibrium. For any subset of agents, the best way to incentivize them is to pay each agent his indifference payment between exerting effort and shirking, while paying zero to others. The scope of recent works focuses on computing the optimal contract [Dütting *et al.*, 2023; Ezra *et al.*, 2023; Dütting *et al.*, 2025b].

In the multi-agent model, payments in the optimal contract may vary significantly across the agents. However, in practice, companies and organizations frequently implement standardized payment schemes. For example, chain stores such as McDonald’s or Starbucks set standard hourly rates to treat every worker equally. A strengthened form of this egalitarianism is found in worker cooperatives, a concept dating back to the English labor movement. In organizations like the UK wholesaler Suma Wholefoods, equal pay is implemented to reflect the view that every worker is a co-owner and all contributions are essential to the business’s success [Ethicalsuperstore, 2025].

Although equal pay brings high cohesion and equality, it is not universally viable. Not only because high-skill workers can command higher wages elsewhere, but also those high performers consider it unfair. A practical trade-off is to adopt the wage ratio, which refers to the gap between the top salaries and the bottom salaries in economy. For example, as a world leader in the co-operative movement, the Mondragon Corporation in Spain uses agreed-upon wage ratios [Institute, 2022]. The ratios range from 1:3 to 1:9 in its different cooperatives. The statutory minimum wage could also be viewed through a similar lens: by having the highest wage fixed, the minimum wage was guaranteed.

In this paper, we study contracts that give agents equitable payments, as considered in [Feng *et al.*, 2024; Castiglioni *et*

* Ankang Sun is the corresponding author.

al., 2025b]. Motivated by wage ratios used in practice, we extend this by allowing agents’ payments to differ by a multiplicative factor. Our goal is to offer a comprehensive understanding of how the degree of non-discrimination impacts the principal’s optimal utility.

1.1 Contribution

We study multi-agent contracts, where a principal delegates a task to multiple agents. Each agent incurs a cost upon exerting effort, and the overall success probability across agents is monotone submodular. The principal incentivizes agents to contribute effort by offering contractual payments, aiming to maximize her utility. Our paper specifically centers on non-discriminatory contracts, which ensure equitable payments to all agents, and we explore their impact on the principal’s optimal utility. The results are summarized in Table 1.

Similar to [Feng *et al.*, 2024], we consider the concept of price of non-discrimination (PoND), which is defined as the supremum ratio between the principal’s optimal utility of unconstrained contracts and that of all non-discriminatory contracts. For any submodular reward function with n agents, we design a non-discriminatory contract that achieves $\Omega(\frac{1}{\log n})$ fraction of the optimal utility without constraints¹, implying an $O(\log n)$ upper bound on PoND. We then prove that the bound cannot be improved by more than an $O(\log \log n)$ factor, even if the success probability function is additive. Thus, the PoND is bounded between $\Omega(\frac{\log n}{\log \log n})$ and $O(\log n)$.

Main Result 1: $\text{PoND} = \tilde{\Theta}(\log n)$.

As n increases, the bound $\Theta(\log n)$ becomes large, which means that optimality and non-discrimination are hardly compatible. We observe that if we are willing to relax the non-discrimination requirement, which is often acceptable in practice, the utility can be substantially increased. To more thoroughly and accurately examine this effect, for $\beta \geq 1$, we extend the notion of non-discrimination to β -approximate non-discriminatory contracts (or β -non-discriminatory contracts), which permit the difference between the highest and lowest payments to be within a factor of β . When $\beta = 1$, all agents receive equal payments; as β approaches infinity, agents are paid according to the optimal unconstrained contract. Accordingly, the price of β -non-discrimination ($\text{PoND}(\beta)$) can be defined in the same manner.

We then provide a comprehensive parametric study on $\text{PoND}(\beta)$, by setting $\beta = n^\delta$, where $\delta \geq 0$. When $\delta = 0$, the problem degenerates to the non-relaxed version and in the following, we assume $\delta > 0$. Our goal is to characterize the relationship between $\text{PoND}(\beta)$ and δ .

Main Result 2: Letting $\beta = n^\delta$,

- for $0 < \delta < 1$, $\max\{2, \frac{1}{\delta + o(1)}\} \leq \text{PoND}(\beta) \leq \lceil \frac{1}{\delta} \rceil + 1$;
- for $\delta = 1$, $1.591 \leq \text{PoND}(\beta) \leq 2$;
- for $\delta > 1$, $1 \leq \text{PoND}(\beta) \leq 1 + o(1)$.

The above result offers an algorithmic framework that assists decision makers in selecting an appropriate level of

non-discrimination, taking into account the potential loss in the principal’s utility associated with non-discriminatory contracts. The bounds are asymptotically tight for both small and large values of δ . For example, if δ is set to $\frac{1}{2}$, then the ratio between the highest and lowest payments is within $\sqrt{n}$, and our contract ensures the principal’s utility to be at least $\frac{1}{3}$ of the optimal utility. If $\delta = \frac{1}{\log n}$, the ratio of highest and lowest payments is within 2, and our contract ensures the principal’s utility to be at least $\frac{1}{\log n + 1}$ of the optimal utility. Besides, the bound is not tight for values of δ close to 1, which remains an open question for future research.

In the previous two sets of results, we focused on the asymptotic bounds as n becomes large. To complement this analysis, we present a case study for two agents, where we derive the exact tight bound on $\text{PoND}(\beta)$ using a more algebraic approach.

Main Result 3: When $n = 2$, $\text{PoND}(\beta) = 1 + \frac{1}{\sqrt{\beta+1}}$.

It is worth noting that non-discriminatory contracts in the same multi-agent model are studied in a simultaneous and independent work [Castiglioni *et al.*, 2025b]. However, their main focus is on quantifying the difference between non-discriminatory contracts and other forms of *fair* contracts, rather than examining how non-discrimination impacts the optimal utility of the principal.

1.2 Related Work

Algorithmic contract design. Following two aforementioned papers in microeconomics by [Holmström, 1979] and [Grossman and Hart, 1992], contract theory has been well explored from an algorithmic perspective. A core problem within is combinatorial contracts, which can be categorized into single-agent multi-action contracts and multi-contracts. [Dütting *et al.*, 2025a] studied the former one where an agent chooses any subset of possible actions, bringing dependencies among agent’s actions. They showed an optimal contract can be computed efficiently if the success probability function is gross substitutes, while computing the optimal contract is NP-hard for submodular functions. As a consequence, a line of works by [Ezra *et al.*, 2023; Dütting *et al.*, 2024; Deo-Campo Vuong *et al.*, 2024; Feldman and Yashin, 2025] also extended the model to a combinatorial setting, with the reward function ranging from gross substitutes and submodular to supermodular. They addressed problems in both tractability and hardness.

Multi-agent contracts. This paper is more closely related to the multi-agent setting pioneered by [Babaioff *et al.*, 2006]. In their setting, a principal incentivizes a team of agents to jointly work on a single project. Each of the n agents chooses to exert effort or not, and optimizing the principal’s utility was defined as the combinatorial agency problem. The landscape was significantly expanded by [Dütting *et al.*, 2023]. In their work, more complicated reward functions were considered, in which the payment to each agent in an incentivized set depends on their marginal contribution to the project. They gave constant-factor approximation algorithms for submodular and XOS reward functions, while there is an $\Omega(\sqrt{n})$ impossibility for subadditive functions with n agents. The most

¹Throughout the paper, $\log$ denotes the base-2 logarithm.

		Upper Bound	Lower Bound
PoND		$O(\log n)$ (Lemma 2)	$\Omega(\frac{\log n}{\log \log n})$ (Lemma 3)
PoND(β) ($\beta = n^\delta$)	$0 < \delta < 1$	$\lceil \frac{1}{\delta} \rceil + 1$ (Lemma 5)	$\max\{2, \frac{1}{\delta + o(1)}\}$ (Lemmas 6 and 7)
	$\delta = 1$	2 (Lemma 5)	1.591 (Lemma 8)
	$\delta > 1$	$1 + o(1)$ (Remark 1)	
	$n = 2$	$1 + \frac{1}{\sqrt{\beta+1}}$ (Theorem 9)	

Table 1: Summary of results for the price of (β -)non-discrimination, where n is the number of agents and β is the payment gap between the linear contract.

recent work by [Ezra *et al.*, 2023] and [Dütting *et al.*, 2025b] showed that no PTAS exists for Submodular functions, and no approximation better than $\Omega(n^{\frac{1}{6}})$ exists.

Another multi-agent model was studied by [Castiglioni *et al.*, 2023] and [Hann-Caruthers and Goel, 2024]. In their setting, each agent completes their own task with their payments conditioned on individual performance. Thus, it is natural that payments to each agent are linked to their own outcomes. Our work, however, is based on the first multi-agent model where a team of agents work on a single project. It is mainly because in practice, organizations like worker cooperatives make everyone work towards a common goal instead of finishing individual tasks. As there is no standard way of judging someone’s outcome, equalizing the payments matters.

Fairness in contracts. [Feng *et al.*, 2024] conducted the study of non-discriminatory contracts. The key difference is that they considered the second multi-agent model as previously mentioned, where agents complete individual tasks instead of a single task. While they have derived logarithmic price of non-discrimination across all their settings, our model introduces the concept of price of β -non-discrimination to achieve constant-factor approximations. Fairness concerns were also investigated by [Castiglioni *et al.*, 2025a], who considered the trade-off between profit-maximizing and fairness when allocating a bunch of tasks to some agents. In a concurrent work with this paper, [Castiglioni *et al.*, 2025b] explore *fair* contracts using a swap-based definition, whose objective is to bound the gap between non-discriminatory contracts and optimal *fair* contracts.

The price of non-discriminatory contracts is also studied in another concurrent and independent work [Feldman *et al.*, 2026]. Their upper bound is stronger than ours: it closes the $\log \log n$ factor gap, holds for XOS success functions, and applies under combinatorial actions. In contrast, in the current paper, we further provide a systematic study of the relaxed notion PoND(β), and give a near-tight bound on any β .

2 Preliminaries

Multi-agent hidden-action setting. For any $k \in \mathbb{N}^+$, let $[k] = \{1, \dots, k\}$. We study the model in which a principal interacts with a set of agents $N = [n] = \{1, \dots, n\}$ to execute a project. Each agent chooses whether to exert

effort or shirk. For each agent i , exerting effort incurs cost $c_i \in \mathbb{R}_{>0}$, while shirking incurs zero cost. The outcome of the project is either “success” or “failure”, and depends on the set of agents who choose to exert effort. The principal observes the outcome of the project but not the actions of agents. Let $f : 2^N \rightarrow [0, 1]$ be the project’s success probability function. When a set of agents S exert effort, the project succeeds with probability $f(S)$. If the project succeeds, the principal derives reward normalized to $r = 1$. Thus, f is also referred to as the (expected) reward function. Otherwise if the project fails, the reward is zero. For any $i \in N$ and $S \subseteq N$, define $f(i|S) := f(S \cup \{i\}) - f(S)$ as the marginal contribution of including i to S . In this paper, f is considered to be monotone and submodular, where monotonicity ensures that for any $S \subseteq T$, $f(S) \leq f(T)$, and submodularity ensures that for any $S \subseteq T$ and any $i \notin T$, $f(i|T) \leq f(i|S)$. For any $S \subseteq N$, $|S|$ denotes its cardinality.

Optimal contract design. Since exerting effort is costly, agents by themselves have no interest in exerting effort, of which the benefit goes to the principal. To incentivize agents, the principal designs a contract that assigns payments to agents based on the outcome of the project. In this paper, we focus on the *linear contract*, an n -dimensional vector $\alpha = (\alpha_1, \dots, \alpha_n)$ with every $\alpha_i \in [0, 1]$. A contract α assigns payment α_i to agent i when the project succeeds, regardless of whether agent i exerts effort or not.

When analyzing contracts, we consider the pure Nash equilibria of the induced game among the agents and adopt the tie-breaking rule where agents exert effort when indifferent. For a contract α , suppose that it incentivizes agents S to exert effort, i.e.,

$$\begin{aligned} \alpha_i f(S) - c_i &\geq \alpha_i f(S \setminus \{i\}) \quad \text{for all } i \in S, \\ \alpha_i f(S) &\geq \alpha_i f(S \cup \{i\}) - c_i \quad \text{for all } i \notin S. \end{aligned}$$

Note that each agent $i \in S$ has expected utility $\alpha_i f(S) - c_i$, while each $i \notin S$ has expected utility $\alpha_i f(S)$. Then for contract α and the equilibrium where agents in S exert effort, the principal’s utility is

$$(1 - \sum_{i=1}^n \alpha_i) f(S).$$

The principal pursues the contract that maximizes her utility. Then for a fixed set of agents S , the contract α that maximizes

the principal’s utility among all contracts in which the set of agents exerting effort is exactly S , is

$$\alpha_i = \frac{c_i}{f(i|S \setminus \{i\})} \quad \text{for all } i \in S,$$

$$\alpha_i = 0 \quad \text{for all } i \notin S.$$

For any $S \subseteq N$, define

$$g(S) := \left(1 - \sum_{i \in S} \frac{c_i}{f(i|S \setminus \{i\})}\right) f(S),$$

the principal’s optimal utility for the contracts where the set of agents exerting effort is S . Therefore, the problem of finding the contract with the maximum utility of the principal is equivalent to $\max_{S \subseteq 2^N} g(S)$, which, however, is known to be NP-hard [Dütting *et al.*, 2023].

Approximately non-discriminatory contracts. We now introduce non-discriminatory (linear) contracts. Informally, a contract is non-discriminatory (ND) if the payments assigned to agents who exert effort are identical. The ND requirement does not apply to the agents who choose to shirk. For a fixed set of agents S , the ND contract α that incentivizes S and maximizes the principal’s utility satisfies

$$\alpha_i = \max_{j \in S} \frac{c_j}{f(j|S \setminus \{j\})} \quad \text{for all } i \in S,$$

$$\alpha_i = 0 \quad \text{for all } i \notin S.$$

For any $S \subseteq N$, define

$$\alpha_S := \max_{j \in S} \frac{c_j}{f(j|S \setminus \{j\})}.$$

The contract that pays α_S to each agent in S and zero to all other agents incentivizes exactly the agents in S and attains the minimum total payment among all contracts that induce effort by S . We also define

$$g_{ND}(S) := (1 - |S|\alpha_S) f(S).$$

The interpretation of $g_{ND}(S)$ is the principal’s optimal utility under the ND contract that induces effort by S . Therefore, in this context, the principal faces the optimization problem of $\max_{S \subseteq 2^N} g_{ND}(S)$. Throughout the paper, we refer to contracts without the ND constraint as *unconstrained* contracts. When incentivizing the same set of agents S , the payments for agents under the optimal ND contract are weakly higher than those under the unconstrained contract. Hence, the non-discrimination requirement reduces the principal’s utility.

In addition to the exact non-discrimination, we are also concerned with approximately non-discriminatory contracts, in which payments assigned to agents who exert effort do not differ substantially. To measure the disparity among agents’ payments, we use a multiplicative factor β with $\beta \geq 1$, standing for the wage ratio. For a β -non-discriminatory (β -ND) contract α , if it incentivizes S to exert effort, then for any $i, j \in S$ with $\alpha_i \leq \alpha_j$, $\frac{\alpha_i}{\beta} \leq \alpha_j$ holds. For a fixed set of agents S , the β -ND contract α that incentivizes S and maximizes the principal’s utility satisfies

$$\alpha_i = \max \left\{ \frac{c_i}{f(i|S \setminus \{i\})}, \frac{1}{\beta} \alpha_S \right\} \quad \text{for all } i \in S,$$

$$\alpha_i = 0 \quad \text{for all } i \notin S.$$

For any $S \subseteq N$ and $\beta \geq 1$, define $g_{ND}(\beta, S)$ as

$$\left(1 - \sum_{i \in S} \max \left\{ \frac{c_i}{f(i|S \setminus \{i\})}, \frac{1}{\beta} \alpha_S \right\} \right) f(S),$$

which stands for the principal’s utility under the optimal β -ND contract in which the set of agents exerting effort is S . Note that $g_{ND}(1, S) = g_{ND}(S)$. For β -ND contracts, the optimization problem for principal is $\max_{S \subseteq 2^N} g_{ND}(\beta, S)$.

The price of non-discrimination. We quantify the optimal utility loss due to the requirement of non-discrimination through the framework of the price of β -non-discrimination, a function with respect to β . Informally, it is the supremum ratio between the principal’s optimal utility of all contracts and the principal’s optimal utility of all β -ND contracts. Denote by $\mathcal{I}$ the set of problem instances. For any instance $I \in \mathcal{I}$, let $\text{OPT}(I)$ be the principal’s optimal utility of all contracts and let $\text{OPT}_{ND}(\beta, I)$ be the principal’s optimal utility of all β -ND contracts.

Definition 1. For any $\beta \geq 1$, the price of β -non-discrimination is

$$\text{PoND}(\beta) := \sup_{I \in \mathcal{I}} \frac{\text{OPT}(I)}{\text{OPT}_{ND}(\beta, I)}.$$

When I and β are clear from the context, we use OPT and OPT_{ND} to refer to $\text{OPT}(I)$ and $\text{OPT}_{ND}(\beta, I)$, respectively. Also, PoND denotes $\text{PoND}(1)$ for simplicity.

3 Logarithmic PoND

In this section, we study the price of (exact) non-discrimination. We show that there always exists a non-discriminatory contract achieving at least $\frac{1}{\log n}$ of the unconstrained optimum, and this is almost tight, as there exist instances where the ratio is at most $O(\frac{\log \log n}{\log n})$.

Theorem 1. For n agents with submodular reward function f , the price of 1-non-discrimination is at least $\Omega(\frac{\log n}{\log \log n})$ and at most $O(\log n)$.

The difficulty for deriving the $O(\log n)$ upper bound lies in identifying the set of agents to be incentivized by the contract. A natural greedy approach that adds agents sequentially, for example in decreasing order of “price to performance ratio²”, is unlikely to work. The rationale is that including an agent with huge payment gap can dramatically reduce the principal’s utility, since under non-discrimination all incentivized agents are equally paid. Skipping that high-reward agent does not necessarily resolve the issue, because the current set of incentivized agents may already be at a threshold where adding any additional agent pushes the total payment above 1.

Instead, we design the contract based on a particular configuration of agents. Begin with the optimal unconstrained contract in which the agents in S^* are incentivized to exert effort. We partition S^* into several groups $G_1, \dots, G_m$, such that (1) the total payment to each group under ND contracts

²For an agent, his price-performance ratio is the ratio between his marginal contribution and his payment in the optimal contract.

	G_1	G_2	...	G_k	...	G_{m-1}	G_m
Agent index	{1}	{2,3}	...	$\{2^{k-1}, \dots, 2^k - 1\}$	...	$\{2^{m-2}, \dots, 2^{m-1} - 1\}$	$\{2^{m-1}, \dots, n\}$
Group size	1	2	...	2^{k-1}	...	2^{m-2}	$< 2^{m-1}$
Optimal payment	$\frac{c_1}{f(1)}$	...	...	$\max_{j \in G_k} \frac{c_j}{f(j G_k \setminus \{j\})}$	...	$\max_{j \in G_{m-1}} \frac{c_j}{f(j G_{m-1} \setminus \{j\})}$	...

Figure 1: Our partition of group $G_1, \dots, G_k$ in Lemma 2. Each group is double the size of the previous one, except for the last group. The optimal payment of group k stands for the standardized payment when incentivizing that single group only.

does not exceed the total payment to S^* in the optimal unconstrained contract, and (2) the sum of the principal's reward over the groups under ND contract is at least the reward under the optimal unconstrained contract.

Lemma 2. *Suppose S^* is the set of agents exerting effort in the optimal unconstrained contract. When f is submodular, there exists a subset of agents $S \subseteq S^*$ such that*

$$g_{ND}(S) = f(S)(1 - |S|\alpha_S) \geq \frac{1}{\lceil \log n \rceil} g(S^*).$$

Proof. Since $|S^*| \leq n$, it suffices to show $g_{ND}(S) \geq \frac{1}{\lceil \log |S^*| \rceil} g(S^*)$. Assume without loss of generality that the optimal unconstrained contract incentivizes all n agents, i.e. $S^* = N$. For each $i \in N$, define $\alpha_{i,N} = \frac{c_i}{f(i | N \setminus \{i\})}$. Reindex the agents according to a descending order of $\alpha_{i,N}$, such that $\alpha_{1,N} \geq \dots \geq \alpha_{n,N}$.

We now partition N into $m = \lceil \log n \rceil$ groups $G_1, \dots, G_m$ as follows: for each $k \in [m-1]$, group G_k contains 2^{k-1} consecutive agents $2^{k-1}, 2^{k-1} + 1, \dots, 2^k - 1$; the remaining agents form group G_m . Informally, group G_m contains the half of N with smaller $\alpha_{i,N}$, and group G_{m-1} contains the half of $N \setminus G_m$ with smaller $\alpha_{i,N}$, and so on. We illustrate this partition in Figure 1.

Consider incentivizing all agents in a single group G_k under ND contracts. The optimal ND contract that incentivizes exactly G_k pays each $i \in G_k$ the amount $\alpha_{G_k} = \max_{j \in G_k} \frac{c_j}{f(j | G_k \setminus \{j\})}$ and pays zero to agents in $N \setminus G_k$. The total payment to agents in G_k equals

$$\begin{aligned}
 |G_k| \cdot \alpha_{G_k} &= |G_k| \cdot \max_{j \in G_k} \frac{c_j}{f(j | G_k \setminus \{j\})} \\
 &\stackrel{(a)}{\leq} |G_k| \cdot \max_{j \in G_k} \frac{c_j}{f(j | N \setminus \{j\})} \\
 &\stackrel{(b)}{\leq} \sum_{i=1}^{|G_k|-1} \frac{c_i}{f(i | N \setminus \{i\})} + \max_{j \in G_k} \frac{c_j}{f(j | N \setminus \{j\})} \\
 &= \sum_{i=1}^{|G_k|} \alpha_{i,N} \leq \sum_{i=1}^n \alpha_{i,N},
 \end{aligned} \tag{1}$$

where (a) holds due to $G_k \subseteq N$ and the submodularity of f and (b) is due to that $\alpha_{i,N}$'s are sorted in descending order and $|G_1| + |G_2| + \dots + |G_{k-1}| \geq |G_k| - 1$.

Inequality (1) implies that for each group G_k , the principal's maximum utility when incentivizing G_k is

pal's maximum utility when incentivizing G_k is

$$g_{ND}(G_k) = f(G_k)(1 - |G_k| \cdot \alpha_{G_k}) \geq f(G_k)(1 - \sum_{i=1}^n \alpha_{i,N}).$$

Summing over $k \in [m]$ yields

$$\begin{aligned}
 \sum_{k=1}^m g_{ND}(G_k) &\geq \left(\sum_{k=1}^m f(G_k) \right) (1 - \sum_{i=1}^n \alpha_{i,N}) \\
 &\geq f(N)(1 - \sum_{i=1}^n \alpha_{i,N}) = \text{OPT},
 \end{aligned}$$

where the second inequality transition is due to the submodularity of f . The above inequality implies that the maximum of $g_{ND}(G_k)$ among all $k \in [m]$ is at least $\frac{1}{m} \text{OPT} = \frac{1}{\lceil \log n \rceil} \text{OPT}$. $\square$

Lemma 2 indicates $\text{PoND} \leq O(\log n)$. For the lower bound, we construct an additive instance with about $\log n$ groups of agents, homogeneous within each group. Incentivizing any group yields the same reward. As shown in this example, the optimal unconstrained contract incentivizes all groups, while any ND contract that yields positive utility for the principal can incentivize at most $O(\frac{\log \log n}{\log n})$ groups.

Lemma 3. *There exists an instance with additive f such that the principal's utility under any ND contract is no greater than a $O(\frac{\log \log n}{\log n})$ fraction of that under the optimal unconstrained contract.*

Proof. Let $n = 2^m - 1$ where m is an integer. In the instance, the n agents are classified into m groups $G_1, G_2, \dots, G_m$, where for any $k \in [m]$, group G_k contains 2^{k-1} agents. Agents within a group share the same cost function and reward function. In particular, for any G_k and any $i \in G_k$, $f(\{i\}) = \frac{1}{m2^{k-1}}$ and $c_i = \frac{1}{Tm2^{k-1}}$, where $T \geq 2$ is a constant number. Thus, to incentivize agent $i \in G_k$, the principal must pay at least $\frac{1}{Tm2^{k-1}}$.

We claim that the optimal utility of the principal under the unconstrained contract is achieved by incentivizing all n agents. If A is a strict subset of N , then there exists an agent $j \in N$ with $j \notin A$. Define $A' := A \cup \{j\}$. One can verify that the difference between the utilities at A' and A equals

$$f(\{j\}) \left(1 - \sum_{i \in A} \alpha_i - \alpha_j \right) - f(A)\alpha_j.$$

Observe for any $i \in N$, $\alpha_i = \frac{f(\{i\})}{T}$. Since $\sum_{i \in A} \alpha_i \leq \frac{1}{T}$ and $f(A) \leq 1$, the difference is at least $f(\{j\})\frac{T-1}{T} - \alpha_j = f(\{j\})\frac{T-2}{T} > 0$. Therefore, $g(S)$ is maximized at $S = N$ with value $1 - \frac{1}{T}$.

Next, we turn to ND contracts and suppose that the optimal ND contract incentivizes only a set $S \subseteq N$. Then in this ND contract, each agent in $N \setminus S$ receives zero payment. Define $L := \min\{k : G_k \cap S \neq \emptyset\}$. Since each agent in S receives the same payment $\frac{1}{Tm2^{L-1}}$ under the optimal ND contract, and agents in the group with a lower index have higher marginal contributions to f , the principal must prioritize agents from lower-index groups. Hence, if S contains some agent(s) in G_k for $k \geq L+1$, then all agents from G_L to G_{k-1} must belong to S .

Assume without loss of generality that $f(S) = \lambda \leq 1$, since $f(N) = 1$. Because every group contributes $\frac{1}{m}$ to the total reward, S must contain at least $t = \lfloor \lambda m \rfloor$ entire groups, that is, $G_L \cup G_{L+1} \cup \dots \cup G_{L+t-1} \subseteq S$. Then the number of agents in S satisfies

$$|S| \geq 2^{L-1} \sum_{j=1}^t 2^{j-1} = 2^{L-1} (2^{\lfloor \lambda m \rfloor} - 1).$$

Since each agent in S receives payment $\frac{1}{Tm2^{L-1}}$, the principal's optimal utility $g_{ND}(S)$ is at most

$$\lambda \left(1 - \frac{2^{L-1} (2^{\lfloor \lambda m \rfloor} - 1)}{Tm2^{L-1}} \right) = \lambda \left(1 - \frac{2^{\lfloor \lambda m \rfloor} - 1}{Tm} \right).$$

Since an ND contract that incentivizes a single group yields positive utility for the principal, the upper bound on $g_{ND}(S)$ must be positive. Then $2^{\lfloor \lambda m \rfloor} - 1 \leq Tm$ holds, and thus, $\lambda \leq \frac{\log(Tm+1)}{m}$. Therefore, it holds that

$$\frac{\text{OPT}}{\text{OPT}_{ND}} > \frac{T-1}{T} \frac{m}{\log(Tm+1)},$$

implying that $\text{PoND} = \Omega(\frac{\log n}{\log \log n})$ as $n = 2^m - 1$. $\square$

Theorem 1 directly follows from Lemmas 2 and 3.

4 Constant-Factor PoND(β)

In the previous bad example, a non-discriminatory contract yields about $O(\frac{1}{\log n})$ of the optimal utility, driven by large disparities in unconstrained payments across agents. A natural question raised is how much the price of non-discrimination improves if the ratio of the maximum to minimum payment among incentivized agents is bounded by a factor β .

In this section, we show that when $\beta = n^\delta$ for some constant $\delta \in (0, 1]$, $\text{PoND}(\beta)$ is bounded by a constant. The case when $\delta > 1$ follows from the remark below:

Remark 1. For $\delta > 1$, one can achieve exact-approximation with sufficiently large n : For agents N incentivized under the optimal contract, the total payment with and without ND differs by at most $\frac{(n-1)\alpha_i}{n} \rightarrow 0$, where α_i is the maximum payment value among N . Thus, we restrict our attention to $\delta \in (0, 1]$.

Theorem 4. Suppose f is submodular. For any $\delta \in (0, 1]$ and sufficiently large n , the price of n^δ -non-discrimination is at most $\lceil \frac{1}{\delta} \rceil + 1$ and at least $\max\{2, \frac{1}{\delta + O(\frac{\log \log n}{\log n})}\}$ when $\delta \in (0, 1)$ and at least 1.591 when $\delta = 1$.

The lemma below establishes the upper bound on the price of n^δ -non-discrimination.

Lemma 5. Suppose that f is submodular and S^* is the set of agents exerting effort in the optimal unconstrained contract. For any $\delta \in (0, 1]$ and sufficiently large n , there exists an ND contract in which agents $S \subseteq S^*$ exert effort and

$$g_{ND}(n^\delta, S) \geq \frac{1}{\lceil \frac{1}{\delta} \rceil + 1} g(S^*).$$

Proof. As our objective is to design an ND contract that incentivizes a subset of agents in S^* , assume without loss of generality that the optimal unconstrained contract incentivizes all n agents as we did in Lemma 2, i.e. $S^* = N$. For any $i \in N$, define $\alpha_{i,N} := \frac{c_i}{f(i|N \setminus \{i\})}$. By the optimality of $g(S^*)$, we have $\alpha_{i,N} < 1$ for all i . We now group agents by their $\alpha_{i,N}$ values.

Let $t = \lceil \frac{1}{\delta} \rceil$ and we partition agents into $t+1$ groups $G_1, G_2, \dots, G_{t+1}$, defined as follows:

$$G_1 = \left\{ i \in N : 0 \leq \alpha_{i,N} < \frac{1}{n} \right\},$$

$$G_2 = \left\{ i \in N : \frac{1}{n} \leq \alpha_{i,N} < \frac{1}{n^{(t-1)\delta}} \right\},$$

and for any $j = 3, \dots, t+1$,

$$G_j = \left\{ i \in N : \frac{1}{n^{(t+2-j)\delta}} \leq \alpha_{i,N} < \frac{1}{n^{(t+1-j)\delta}} \right\}.$$

Note that group G_2 is well defined because $(t-1)\delta < 1$. Starting from the right boundary of G_2 , repeatedly extend to the right by a factor of $\frac{1}{\delta}$ to obtain groups $G_3, \dots, G_{t+1}$.

We then consider β -ND contracts that incentivize each group G_k separately and show that the sum of the principal's utilities from these contracts is at least the optimal utility under the unconstrained contract. For any agent i , suppose that $i \in G_k$. Define

$$\hat{\alpha}_i := \max \left\{ \frac{c_i}{f(i|G_k \setminus \{i\})}, \frac{1}{n^\delta} \max_{j \in G_k} \frac{c_j}{f(j|G_k \setminus \{j\})} \right\}$$

the minimum payment to i under the n^δ -ND contract that induces effort by exactly G_k .

First, consider group G_k with $k \geq 2$. Since f is submodular, then $\sum_{i \in G_k} \hat{\alpha}_i$ is at most

$$\sum_{i \in G_k} \max \left\{ \frac{c_i}{f(i|N \setminus \{i\})}, \frac{1}{n^\delta} \max_{j \in G_k} \frac{c_j}{f(j|N \setminus \{j\})} \right\},$$

which equals $\sum_{i \in G_k} \alpha_{i,N}$ as $\alpha_{i,N} \geq \frac{1}{n^\delta} \max_{j \in G_k} \alpha_{j,N}$.

For group G_1 , since $\hat{\alpha}_i \leq \alpha_{i,N} < \frac{1}{n}$ for all $i \in G_1$, we have

$$\sum_{i \in G_1} \hat{\alpha}_i < \frac{|G_1|}{n^{(1+\delta)}} + \sum_{i \in G_1} \alpha_{i,N}.$$




Figure 2: Illustration of our example in Lemma 6. The blue part (A) gets the same payment as in the optimal unconstrained contract, while agents in green part (B) receive a uniform payment, which equals the highest payment to agents in A divided by n^δ .

Since $|G_1| \leq n$ and n is sufficiently large, the above inequality implies $\sum_{i \in G_1} \hat{\alpha}_i < \sum_{i \in G_1} \alpha_{i,N} + o(1)$. Note that setting payment $\hat{\alpha}_i$ to every $i \in G_k$ and zero to others yields the optimal n^δ -ND contract that induces effort exactly by G_k . Under this contract, the principal has utility

$$f(G_k) \left(1 - \sum_{i \in G_k} \hat{\alpha}_i \right) = g_{ND}(n^\delta, G_k).$$

We design such a contract for each group G_k . Then the principal's total utility from $t+1$ separate contracts is at least

$$\begin{aligned} & \sum_{k=1}^{t+1} f(G_k) \left(1 - \sum_{i \in G_k} \alpha_{i,N} \right) - f(G_1) \cdot o(1) \\ & \geq f(G_1) \left(1 - \sum_{i \in N} \alpha_{i,N} \right) + \sum_{k=2}^{t+1} f(G_k) \left(1 - \sum_{i \in N} \alpha_{i,N} \right) \\ & \geq f(N) \left(1 - \sum_{i \in N} \alpha_{i,N} \right) = \text{OPT}, \end{aligned}$$

where the first inequality transition is due to $G_k \subseteq N$ for all k and the second inequality uses submodularity of f . By the above inequality, there exists a group G_{k^*} such that the n^δ -ND contract with payment $\hat{\alpha}_i$ to every $i \in G_{k^*}$ and 0 to all other agents yields utility for the principal of at least $\frac{1}{t+1} \text{OPT} = \frac{1}{\lfloor \frac{1}{\delta} \rfloor + 1} \text{OPT}$. $\square$

We complement the upper bound with a counter example, in which the optimal n^δ -ND contract achieves at most $\delta + o(1)$ fraction of the unconstrained optimum.

Lemma 6. *For any $\delta \in (0, 1)$, there exists an instance where the principal's optimal utility under any n^δ -ND contract is at most a $(\delta + O(\frac{\log \log n}{\log n}))$ fraction of the optimal utility under unconstrained contracts.*

We outline the proof of Lemma 6 and the proof could be found in the full version of this paper [Ding et al., 2026]. To sketch, we consider the very same instance as that in Lemma 3, where the agents belonging to lower-indexed groups have higher success probability and payment threshold, and vice versa. Based on our definition of β -ND, we partition all

agents into A and B according to a “pivot” group k_0 , as illustrated in Figure 2. Agents in groups with index smaller than or equal to k_0 are paid the original amount as in the optimal unconstrained contract, while agents in groups with higher index are paid a uniform amount. We upper-bound OPT_{ND} by incentivizing as many agents from A, plus some consecutive groups from B.

We also derive another lower bound that improves the bound in Lemma 6 in some cases, e.g., when $\delta \geq \frac{1}{2}$.

Lemma 7. *For any $\delta \in (0, 1)$ and arbitrarily small $\epsilon > 0$, there exists an instance where the principal's optimal utility under n^δ -ND contracts is at most a $(\frac{1}{2} + \epsilon)$ fraction of that under unconstrained contracts.*

The details of the construction are given in our full version [Ding et al., 2026]. The idea is as follows: the upper bound shown in Lemma 5 requires payment disparities. We exploit such disparities by creating an agent with high contribution and payment (A) and a group of “free labor” with very low contribution and payment (B). In order to incentivize the powerful agent in A, the principal is forced to exclude all the agents in B.

To conclude, we provide an lower bound for $\delta = 1$.

Lemma 8. *There exists an instance where the principal's optimal utility under n -ND contracts is no greater than 0.629 fraction of that under unconstrained contracts.*

Basically, the construction is analogous to that of Lemma 7, while we can additionally incentivize a portion of agents in B. The discussion could be found in the full version [Ding et al., 2026].

5 Tight Results for Two Agents

Our results in the previous section are grounded in large n . In this section, we shift our focus to the case of two agents, in which we are able to identify the optimal β -ND contract, and establish tight $\text{PoND}(\beta)$ for all $\beta \geq 1$.

Theorem 9. *For any $\beta \geq 1$, when there are two agents and the reward function f is submodular, the $\text{PoND}(\beta)$ equals $(1 + \frac{1}{\sqrt{\beta+1}})$, and this bound is tight.*

The whole proof of this theorem is given in the full version [Ding et al., 2026]. The $\text{PoND}(\beta)$ is characterized by enumerating all different ways to incentivize agent(s).

6 Conclusion

This paper investigates how varying degrees of non-discrimination influence the optimal revenue in multi-agent contracts. We provide asymptotic bounds on the price of β -non-discrimination, characterizing the tradeoff between non-discrimination and revenue. Our bounds are nearly tight for both the strict and relaxed cases when β is either small or large. However, determining the exact bound when β approaches n remains an open question for future research.

Our results uncover several promising directions for further study. For example, it would be valuable to extend our framework to other classes of reward functions, such as fractionally subadditive (XOS) and subadditive functions. It is also interesting to investigate alternative objectives, such as social welfare, social surplus, and various fairness criteria.

Acknowledgments

We thank Yiding Feng and Yingkai Li for valuable discussions, and the anonymous reviewers for their helpful comments. Bo Li is supported in part by the Hong Kong SAR Research Grants Council (Grant No. PolyU 15224823) and Guangdong Basic and Applied Basic Research Foundation (Grant No. 2024A1515011524).

References

- [Babaioff *et al.*, 2006] Moshe Babaioff, Michal Feldman, and Noam Nisan. Combinatorial agency. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 18–28, 2006.
- [Castiglioni *et al.*, 2023] Matteo Castiglioni, Alberto Marchesi, and Nicola Gatti. Multi-agent contract design: How to commission multiple agents with individual outcomes. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 412–448, 2023.
- [Castiglioni *et al.*, 2025a] Matteo Castiglioni, Junjie Chen, and Yingkai Li. Fair contracts. *CoRR*, abs/2507.11214, 2025.
- [Castiglioni *et al.*, 2025b] Matteo Castiglioni, Junjie Chen, and Yingkai Li. Fair team contracts. *CoRR*, abs/2512.19388, 2025.
- [Dai Li *et al.*, 2021] Wanyi Dai Li, Nicole Immorlica, and Brendan Lucier. Contract design for afforestation programs. In *International Conference on Web and Internet Economics*, pages 113–130. Springer, 2021.
- [Deo-Campo Vuong *et al.*, 2024] Ramiro Deo-Campo Vuong, Shaddin Dughmi, Neel Patel, and Aditya Prasad. On supermodular contracts and dense subgraphs. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 109–132. SIAM, 2024.
- [Department, 2021] Statista Research Department. www.statista.com/statistics/950920/global-instagram-influencer-marketing-spending, 2021.
- [Ding *et al.*, 2026] Ke Ding, Bo Li, and Ankang Sun. Multi-agent non-discriminatory contracts. *CoRR*, abs/2601.16835, 2026.
- [Dütting *et al.*, 2023] Paul Dütting, Tomer Ezra, Michal Feldman, and Thomas Kesselheim. Multi-agent contracts. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1311–1324, 2023.
- [Dutting *et al.*, 2024] Paul Dutting, Michal Feldman, and Yoav Gal Tzur. Combinatorial contracts beyond gross substitutes. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 92–108. SIAM, 2024.
- [Dütting *et al.*, 2025a] Paul Dütting, Tomer Ezra, Michal Feldman, and Thomas Kesselheim. Combinatorial contracts. *SIAM Journal on Computing*, 54(6):1427–1455, 2025.
- [Dütting *et al.*, 2025b] Paul Dütting, Tomer Ezra, Michal Feldman, and Thomas Kesselheim. Multi-agent combinatorial contracts. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1857–1891. SIAM, 2025.
- [Ethicalsuperstore, 2025] Ethicalsuperstore. www.ethicalsuperstore.com/blog/brand-ethics/suma-wholefoods-ethical-sustainable-and-worker-owned, 2025.
- [Ezra *et al.*, 2023] Tomer Ezra, Michal Feldman, and Maya Schlesinger. On the (in) approximability of combinatorial contracts. *arXiv preprint arXiv:2311.18425*, 2023.
- [Feldman and Yashin, 2025] Michal Feldman and Liat Yashin. Ultra-efficient contracts: Breaking the substitutes barrier in combinatorial contracts. *arXiv preprint arXiv:2506.18008*, 2025.
- [Feldman *et al.*, 2026] Michal Feldman, Yoav Gal Tzur, Tomasz Ponitka, and Maya Schlesinger. Equal-pay contracts. *CoRR*, abs/2601.15478, 2026.
- [Feng *et al.*, 2024] Yiding Feng, Mengfan Ma, and Mingyu Xiao. Price of non-discrimination in public combinatorial contracts. In *Web and Internet Economics - 20th International Conference, WINE*, pages 25–41. Springer, 2024.
- [Grossman and Hart, 1992] Sanford J Grossman and Oliver D Hart. An analysis of the principal-agent problem. In *Foundations of insurance economics: Readings in economics and finance*, pages 302–340. Springer, 1992.
- [Hann-Caruthers and Goel, 2024] Wade Hann-Caruthers and Sumit Goel. Optimality of weighted contracts for multi-agent contract design with a budget. In *Proceedings of the 25th ACM Conference on Economics and Computation*, page 1295. ACM, 2024.
- [Ho *et al.*, 2014] Chien-Ju Ho, Aleksandrs Slivkins, and Jennifer Wortman Vaughan. Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 359–376, 2014.
- [Holmström, 1979] Bengt Holmström. Moral hazard and observability. *The Bell journal of economics*, pages 74–91, 1979.
- [Holmström, 1982] Bengt Holmström. Moral hazard in teams. *The Bell journal of economics*, pages 324–340, 1982.
- [Institute, 2022] Australian Cardijn Institute. australiancardijninstitute.org/mondragon-limits-maximum-salaries-to-6x/, 2022.