Simple yet Effective Budget-Feasible Procurement Auctions for Submodular Welfare Maximization

Junxiang Zhang[1] **Qingchun HE**[1] **Xinhui Lu**[2] **Yaxin Hu**[2]
Feng Li[3] **Jing Tang**[1] **Kai Han**[2]*

[1] The Hong Kong University of Science and Technology (Guangzhou), P.R.China
[2] School of Computing and Artificial Intelligence, Shanghai University of Finance and Economics
Ministry of Education's Key Laboratory of Interdisciplinary Research of Computation and Economics
Shanghai, 200433, PR China
[3] School of Computer Science and Technology, Shandong University, P.R.China

Abstract

We study budget-feasible procurement auctions for social-welfare maximization under submodular valuations. In this setting, a buyer with a limited payment budget seeks to procure goods or services from strategic sellers, with the objective of maximizing the buyer's value for the selected sellers minus their true costs. We propose a simple yet effective single-round clock auction in which each seller receives at most one price offer. The mechanism satisfies desirable economic properties, including truthfulness, individual rationality, budget feasibility, and non-negative auctioneer surplus. Moreover, it achieves approximation ratios of 8.52 for monotone submodular valuations, 23.2 in expectation for non-monotone submodular valuations, and 24.52 deterministically for non-monotone submodular valuations, using only a linear number of value-oracle queries. These guarantees improve over the recent work of Cui et al. (ICML 2026) on the same problem in approximation ratio, value-oracle query complexity, and number of pricing rounds. Experiments on influence maximization in social networks and crowdsourcing further demonstrate the effectiveness and efficiency of our approach.

1 Introduction

Procurement auctions, also known as reverse auctions, constitute a fundamental framework in auction theory, where a buyer seeks to procure goods or services from a set N of sellers. Each seller $v \in N$ has a privately known cost $c(v)$ for providing the good or service and may act strategically to maximize their utility. A central challenge in mechanism design is therefore to design allocation and payment rules that incentivize truthful behavior while optimizing a relevant objective. Two primary objectives that have received significant attention in the literature are maximizing the buyer's valuation and maximizing social welfare [40].

In many practical settings, the buyer faces a strict budget constraint B, which limits the set of selected sellers and the attainable value. Motivated by this, Singer [42] initiated the study of budget-feasible mechanisms, aiming to maximize the buyer's valuation under the budget constraint. The valuation is modeled by a set function $f(S)$ over the selected sellers $S \subseteq N$. A particularly important and widely studied class of valuation functions is the class of *submodular* functions, which encode the diminishing marginal returns property [7].

Since the seminal work of Singer [42], budget-feasible procurement auctions have attracted significant attention (e.g.,[4–6, 10, 21–23, 26]) due to their broad applications in various domains, including

*Kai Han (Email: hankai@mail.shufe.edu.cn) is the corresponding author and has made the principal contributions to this paper.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

Table 1: Comparison with the recent work of [13] on theoretical performance guarantees. Here, "monotone" and "non-monotone" indicate the type of submodular valuation, "det." and "rand." denote deterministic and randomized mechanisms, respectively, and OPT denotes the optimal welfare.

Work	**Setting**	**Ratio (μ)**	**Query Complexity**	**Pricing Rounds**
Cui et al. [13] (ICML'26)	Monotone, det.	11.4	$\mathcal{O}(n\log(\mathrm{OPT}/\epsilon))$	$\mathcal{O}(\log(\mathrm{OPT}/\epsilon))$
	Non-monotone, det.	30.49	$\mathcal{O}(n\log(\mathrm{OPT}/\epsilon))$	$\mathcal{O}(\log(\mathrm{OPT}/\epsilon))$
This work	Monotone, det.	$\mathbf{8.52}$	$\mathcal{O}(\mathbf{n})$	$\mathbf{1}$
	Non-monotone, rand.	$\mathbf{23.2}$	$\mathcal{O}(\mathbf{n})$	$\mathbf{1}$
	Non-monotone, det.	$\mathbf{24.52}$	$\mathcal{O}(\mathbf{n})$	$\mathbf{1}$

crowdsourcing [2, 33, 44, 47], influence maximization [43], recommender systems [14], federated learning [45, 46], and collective decision-making [32]. However, nearly all existing studies in this line of work concentrate on maximizing the buyer's valuation, while neglecting another central objective in auction design, namely social welfare maximization. In procurement auctions, social welfare is defined as the buyer's valuation for the selected set of sellers minus the sellers' true costs, i.e., $f(S) - \sum_{v\in S} c(v)$, and therefore captures the net economic value generated by the allocation. Under a payment budget, social welfare maximization is crucial for allocative efficiency, ensuring that limited resources are allocated to sellers with the highest net social value.

However, due to the different objective of maximizing the buyer's valuation, i.e., $f(S)$, traditional techniques in budget-feasible mechanism design cannot be applied to maximizing the social welfare , i.e., $f(S) - \sum_{v\in S} c(v)$. This problem is particularly interesting when the valuation function $f(\cdot)$ is submodular, since finding a good solution requires submodular optimization beyond non-negativity; in particular, $f(S) - \sum_{v\in S} c(v)$ may be negative for an improperly chosen set S. In fact, even when the sellers' costs are known and truthfulness is not a concern, maximizing $f(S) - \sum_{v\in S} c(v)$ falls within the scope of *regularized submodular optimization* (RSO) and remains highly non-trivial even under a simple cardinality constraint $|S| \leq k$ [24]. However, existing techniques for RSO usually do not address truthfulness. Therefore, a clear gap remains between existing techniques in auction design and submodular optimization, on the one hand, and the design of budget-feasible procurement auctions for social welfare maximization, on the other hand.

To the best of our knowledge, there is only one very recent independent study [13] addressing the aforementioned gap. The authors propose deterministic, budget-feasible procurement auctions for social welfare maximization under submodular valuations. The winner set S returned by their mechanism satisfies the additive–multiplicative approximation guarantee $f(S) - \sum_{v\in S} c(v) \geq f(S^*)/\mu - \sum_{v\in S^*} c(v) - \lambda(\epsilon)$, where S^* denotes an optimal solution and $\lambda(\epsilon)$ is an additive loss increasing with ϵ. In particular, they obtain $\mu \approx 11.4$ for monotone submodular valuations and $\mu \approx 30.49$ for non-monotone submodular valuations, while requiring $\mathcal{O}(n\log(\mathrm{OPT}/\epsilon))$ value-oracle queries to $f(\cdot)$ in each case, where OPT denotes the optimal social welfare.

However, the theoretical performance guarantees in Cui et al. [13] are subject to several important limitations. First, their approximation ratios (i.e., μ) are relatively large for both monotone and non-monotone submodular valuations. Second, their value-oracle query complexity, $\mathcal{O}(n\log(\mathrm{OPT}/\epsilon))$, depends on the magnitude of OPT, which may be exponentially large in the input size. Third, their mechanism relies on a multi-round clock-auction framework, requiring $\mathcal{O}(\log(\mathrm{OPT}/\epsilon))$ rounds of price offers and bidder responses. In practice, such repeated interactions increase communication costs and prolong the execution of the mechanism, thereby imposing additional participation burdens on sellers and potentially enlarging the scope for strategic behavior [3, 8, 11, 31, 34].

1.1 Our Contributions

In this paper, we propose a simple, effective, and general budget-feasible procurement auction for social welfare maximization under both monotone and non-monotone submodular valuations. In addition to satisfying desirable economic properties, including truthfulness, individual rationality, and non-negative auctioneer surplus, our auction overcomes the limitations of existing work [13] by providing the following performance guarantees:

- Our auction can be implemented either deterministically or randomly, and applies to both monotone and non-monotone submodular valuations. In all cases, the winner set S returned by our auction achieves improved approximation ratios. Specifically, it guarantees $f(S) - \sum_{v \in S} c(v) \geq f(S^*)/\mu - \sum_{v \in S^*} c(v)$, where $\mu \approx 8.52$ is the deterministic approximation ratio for monotone submodular valuations, $\mu \approx 23.2$ is the expected approximation ratio for non-monotone submodular valuations, and $\mu \approx 24.52$ is the deterministic approximation ratio for non-monotone submodular valuations.
- In all cases, our auction uses a linear number of value-oracle queries to $f(\cdot)$. More precisely, for monotone valuation functions, it requires at most $n + 2$ value-oracle queries to $f(\cdot)$. For non-monotone submodular valuations, the number of incurred value-oracle queries is at most $n + 2$ for the randomized implementation and at most $2n + 2$ for the deterministic implementation.
- Our auction requires only one-round pricing, and each seller is offered at most one price, thereby avoiding the drawbacks of multi-round pricing mentioned above.

For clarity, Table 1 compares the theoretical performance guarantees of our work with those of [13]. We also conduct extensive experiments on applications to influence maximization in social networks and crowdsourcing, and the results corroborate the superiority of our approach.

1.2 Technical Overview and Novelty

The auction of Cui et al. [13] essentially adapts the framework of Balkanski et al. [4] for budget-feasible maximization of the buyer's value. In both works, seller prices are determined by guessing the optimal welfare OPT, where the guesses increase geometrically over multiple rounds, resulting in multi-round clock auctions. The underlying intuition is that the problem becomes significantly easier once OPT is known, and the guessing procedure is reminiscent of the doubling trick commonly used in online algorithms when the target benchmark is unknown [41]. However, this repeated guessing procedure is largely mechanical: it introduces additional computational and communication overhead, while making only limited use of the pricing history to improve performance. In particular, the $\log(\mathrm{OPT}/\epsilon)$ factor in both their query complexity and their number of pricing rounds arises from discretizing the interval $[\epsilon, \Theta(\mathrm{OPT})]$ for the guessing process.

Compared with [13], our auction adopts a fundamentally different framework and avoids guessing OPT. It processes sellers sequentially in an arbitrary order, offering each seller v at most one price and adding them to a candidate winner set if they accept. The price is set by a novel rule that balances v's "blunted marginal surplus" against their "opportunity cost": the former measures the welfare gain from considering v, while the latter captures the budget consumed and the lost opportunity to accept better future sellers. This opportunity cost grows exponentially with the payment to v and also increases with the surplus of the current candidate set (i.e., the difference between its value and total payment), making the auction more conservative toward less cost-effective sellers, especially when a large surplus has already been accumulated. Finally, the auction selects a high-quality subset of the candidates as the winners. Despite its simpler design, this economically motivated auction achieves both better approximation guarantees and lower query complexities than Cui et al. [13].

2 Related Work

Budget-feasible mechanisms have been extensively studied for various types of valuation functions, such as additive function [20], submodular function [1, 10, 21–23, 26, 42], XOS and subadditive functions [6, 37, 38]. However, all these studies concentrate on maximizing the buyer's value, and hence cannot be applied to welfare maximization due to the different optimization goals.

There is also a line of work on regularized submodular optimization (RSO), which aims to maximize $f(S) - c(S)$ when element costs are known. Harshaw et al. [24] argue that, since $f(\cdot) - c(\cdot)$ is submodular but may take negative values, an appropriate approximation-guarantee is of the form $f(S) - \sum_{v \in S} c(v) \geq f(S^*)/\mu - \sum_{v \in S^*} c(v)$, and propose an offline distorted greedy framework. Jin et al. [27] develop the offline ROI-greedy algorithm for unconstrained RSO and apply it to profit maximization in social-network influence propagation. Other studies design streaming and distributed algorithms for RSO under various constraints [12, 28, 39]. However, these works do not address the truthfulness issue in the procurement-auction setting.

To the best of our knowledge, the closest studies to our work are [13] (reviewed above) and Deng et al. [15]. In particular, Deng et al. [15] are the first to address truthfulness in RSO, proposing procurement auctions for social welfare maximization with monotone submodular valuations under cardinality constraints or without constraints. However, because their payments required for truthfulness are not subject to any budget constraint, their methods are not applicable to our budget-feasible setting.

3 Preliminaries

We consider a *Budget-Feasible Procurement Auction (BFPA)* consisting of a buyer with budget B and a set N of n sellers, each with a private cost $c(v) \geq 0$. The buyer has a nonnegative submodular valuation function $f : 2^N \to \mathbb{R}_{\geq 0}$ satisfying the diminishing-returns property:

$$\forall X \subseteq Y \subseteq N,\ \forall z \notin Y : f(X \cup \{z\}) - f(X) \geq f(Y \cup \{z\}) - f(Y).$$

Equivalently, f satisfies $\forall X, Y \subseteq N : f(X) + f(Y) \geq f(X \cup Y) + f(X \cap Y)$. Moreover, f may be monotone (i.e., $\forall X \subseteq Y : f(X) \leq f(Y)$) or non-monotone. Following [13], we assume that f is normalized, i.e., $f(\emptyset) = 0$. In a BFPA, the buyer aims to select a set $S \subseteq N$ of sellers and pay each $v \in S$, called a *winner*, an amount $p(v)$ to compensate for their cost, such that the buyer's surplus, i.e., $f(S) - \sum_{v \in S} p(v)$, is maximized under the budget constraint $\sum_{v \in S} p(v) \leq B$. Meanwhile, each seller $v \in N$ aims to maximize their own utility, defined as $p(v) - c(v)$ if $v \in S$, and 0 otherwise. The resulting social welfare is the total utility of all sellers plus the buyer's surplus, i.e., $f(S) - \sum_{v \in S} c(v)$. Since the sellers' costs are private information, they may behave strategically to maximize their utilities. Therefore, to ensure desirable incentives and voluntary participation in a BFPA, we require: (1) *truthfulness*, i.e., each seller maximizes their utility by behaving truthfully, regardless of the strategies adopted by other sellers; (2) *individual rationality (IR)*, i.e., each truthful seller obtains a non-negative utility; and (3) *non-negative auctioneer surplus (NAS)*, i.e., the buyer obtains a non-negative surplus. The performance of a BFPA is measured against the strongest benchmark, namely $f(S^*) - \sum_{v \in S^*} c(v)$, where

$$S^* \in \arg\max\big\{f(T) - \sum\nolimits_{v \in T} c(v) : T \subseteq N,\ \sum\nolimits_{v \in T} c(v) \leq B\big\}. \tag{1}$$

In other words, S^* is the optimal winner set even when all sellers' true costs are publicly known. Note that finding S^* is an NP-hard Regularized Submodular Optimization (RSO) problem, even when the budget constraint degenerates to a cardinality constraint [24]. Following [13] and the literature on RSO (e.g., [12, 24, 28, 39]), we say that S achieves an approximation ratio of μ if S satisfies

$$f(S) - \sum\nolimits_{v \in S} c(v) \geq f(S^*)/\mu - \sum\nolimits_{v \in S^*} c(v), \tag{2}$$

which implies that a smaller μ corresponds to better performance. Even a feasible set with positive value can have negative welfare after costs are deducted. The regularized benchmark discounts only the optimal set's value. We also require a BFPA to be computationally efficient, under the assumption that a *value oracle* is available to return $f(X)$ for any $X \subseteq N$. Since value-oracle queries can be much more time-consuming than other operations, we follow the standard convention in the literature, including submodular optimization (e.g., [9, 12, 15–17, 19, 28, 39]) and budget-feasible mechanism design (e.g., [4, 5, 13, 21–23]), and measure the time complexity of a BFPA by the number of value-oracle queries it makes.

We adopt the fundamental clock auction framework in game theory [35], as briefly reviewed below. In a (descending) clock auction, sellers receive non-increasing price offers over multiple rounds and may reject an offer and permanently exit at any round. The auction may stop after any round, after which the final winners are selected from the remaining active sellers, and each winner is paid the last price they received. It is well known that clock auctions naturally ensure truthfulness and IR [4, 35].

For ease of exposition, for any $X \subseteq N$, let $c(X)$ and $p(X)$ denote $\sum_{v \in X} c(v)$ and $\sum_{v \in X} p(v)$, respectively; let $f(X \mid Y)$ denote $f(X \cup Y) - f(Y)$ for any $X, Y \subseteq N$; for any $v \in N$, let $f(v)$ denote $f(\{v\})$, and let $f(v \mid Y)$ denote $f(\{v\} \mid Y)$; for any $k \in \mathbb{N}^+$, let $[k]$ denote $\{1, \ldots, k\}$.

4 A General Clock Auction for Budget-Feasible Procurement

In this section, we describe the design of our budget-feasible procurement auction for social welfare maximization, as specified in Algorithm 1. Algorithm 1 is essentially a clock auction with only

Algorithm 1 A General Clock Auction for Budget-Feasible Procurement

Require: A set N of n sellers, a submodular valuation function $f(\cdot)$, a budget $B > 0$, parameters $\alpha > 1$, $\beta > 0$, $l \in \{1, 2\}$, $\theta \in (0, 1]$

Ensure: A set $S \subseteq N$ with $\sum_{u \in S} p(u) \leq B$ and payments $\{p(v)\}_{v \in S}$.

1: **Stage 1: One-round pricing:**
2: Initialize the candidate winner sets $A_i \leftarrow \emptyset$ for each $i \in [l]$;
3: **for** each seller $v \in N$ (processed in an arbitrary order) **do**
4: Let $j \in \arg\max_{i \in [l]} f(v \mid A_i)$;
5: **if** $f(v \mid A_j) > 0$ **then**
6: Compute the price $\hat{p}(v)$ to be the value satisfying

$$f(v \mid A_j) - \alpha \cdot \hat{p}(v) = (f(A_j) - p(A_j)) \left[e^{\beta \hat{p}(v)/B} - 1 \right] \tag{3}$$

7: Let $p(v) \leftarrow \min\{B, \hat{p}(v)\}$ and offer the price $p(v)$ to seller v;
8: **if** seller v accepts the price (i.e., $c(v) \leq p(v)$) **then**
9: **with probability** θ **do** $A_j \leftarrow A_j \cup \{v\}$
10: **end if**
11: **end if**
12: **end for**
13: **Stage 2: Selecting the final winner set:**
14: Let $A^* \in \arg\max_{X \in \{A_i : i \in [l]\}} [f(X) - p(X)]$; Let $S \leftarrow A^*$;
15: **if** $p(A^*) > B$ **then**
16: Let H be the smallest subset of A^* consisting of latest-added sellers such that $p(H) > B$;
17: Let w be the earliest-added seller in H; Let $S \in \arg\max_{X \in \{\{w\}, H \setminus \{w\}\}} [f(X) - p(X)]$;
18: **end if**
19: **return** S and the payment vector $\{p(v)\}_{v \in S}$.

one round of pricing, thereby avoiding the shortcomings of multi-round clock auctions mentioned in Section 1. Its framework is general enough to accommodate both monotone and non-monotone submodular valuation functions, and is governed by two input parameters, $l \in \{1, 2\}$ and $\theta \in (0, 1]$. In particular, for monotone valuation functions, it reduces to a deterministic auction when $l = 1$ and $\theta = 1$; for non-monotone submodular valuation functions, it admits both a deterministic implementation when $l = 2$ and $\theta = 1$, and a randomized implementation when $l = 1$ and $\theta \in (0, 1)$. A more detailed analysis of the effects of the parameters l, θ, α, and β in Algorithm 1 on its performance will be presented in Section 5.

In the one-round pricing stage (i.e., Stage 1), Algorithm 1 first constructs l candidate winner sets A_i for $i \in [l]$, and then processes all sellers in an arbitrary fixed order independent of sellers' private costs and actions. Each seller $v \in N$ is greedily assigned to the set A_j that maximizes the marginal value $f(v \mid A_j)$. For every seller with $f(v \mid A_j) > 0$, we then define $\hat{p}(v)$ as the unique solution to Eqn. (3). This solution is well defined: the difference between the left-hand side and the right-hand side of Eqn. (3), viewed as a function of $\hat{p}(v)$, is continuous and strictly decreasing on $[0, f(v \mid A_j)/\alpha]$, and it is positive at the left endpoint and nonpositive at the right endpoint. Hence, Eqn. (3) admits a unique solution in $[0, f(v \mid A_j)/\alpha]$. A complete proof is deferred to Lemma A.2 in Appendix A. After computing $\hat{p}(v)$, the mechanism offers seller v the price $p(v) = \min\{B, \hat{p}(v)\}$. If the offer is accepted, then v is added to A_j with probability θ (Lines 7–9).

After all sellers have been processed, the algorithm enters Stage 2 and selects a set $A^* \in \{A_i : i \in [l]\}$ with the maximum surplus (Line 14). The set A^* is returned as the final winner set if $p(A^*) \leq B$. Otherwise, to satisfy the budget constraint, the algorithm removes from A^* some sellers that were added earliest (i.e., sellers in $A^* \setminus H$), and selects the final winners from the remaining sellers in H (Lines 16–17).

From the pricing stage of Algorithm 1, we obtain the following observation.

Observation 4.1. *Let $A_{i,v}$ denote the state of A_i for each $i \in [l]$ immediately before seller $v \in N$ is processed. We have $\forall i \in [l], \forall v \in A_i : f(v \mid A_{i,v}) - p(v) \geq f(v \mid A_{i,v}) - \alpha\, p(v) \geq 0$. Therefore, throughout the execution of Algorithm 1, both $f(A_i) - \alpha p(A_i)$ and $f(A_i) - p(A_i)$ remain non-negative and non-decreasing for every $i \in [l]$.*

Proof. Both surpluses start at zero. Suppose they are nonnegative before seller v is processed. If $\Delta = f(v \mid A_j) > 0$, the current surplus $f(A_j) - p(A_j)$ is nonnegative by induction. The left side of Eqn. (3) decreases strictly with the value of $\hat{p}(v)$, while the right side is nondecreasing. Their difference is positive when $\hat{p}(v) = 0$ and nonpositive when $\hat{p}(v) = \Delta/\alpha$. Hence $0 \le p(v) \le \hat{p}(v) \le \Delta/\alpha$. Since $\alpha > 1$, the increments upon insertion satisfy $\Delta - p(v) \ge \Delta - \alpha p(v) \ge 0$. All other sets are unchanged, as are all sets on a step without an insertion. This proves both assertions by induction over the sellers. □

Since Algorithm 1 is a clock auction, it is truthful and individually rational by construction. We also prove that Algorithm 1 achieves NAS and budget-feasibility in Appendix B. So we get the following lemma:

Lemma 4.1. *Algorithm 1 achieves truthfulness, IR, NAS, and budget-feasibility.*

Moreover, the pricing rule in Line 6 admits an economic interpretation, as explained below. In fact, it incorporates not only the direct payment cost of selecting a seller, but also the "opportunity cost" of consuming the limited budget. In particular, the term $e^{\beta\hat{p}(v)/B} - 1$ maps the relative budget consumption $\hat{p}(v)/B$ into an increasing and convex penalty. Thus, small payments are penalized approximately linearly, whereas large payments are penalized disproportionately, reflecting the fact that spending a significant fraction of the budget reduces future feasible choices more severely. Multiplying this term by the current surplus $f(A_j) - p(A_j)$ further makes the required marginal gain depend on the quality of the candidate winner set already accumulated. Writing $C = f(A_j) - p(A_j)$ and $\Delta = f(v \mid A_j)$, Eqn. (3) can be read as $C + \Delta - \alpha\hat{p}(v) = Ce^{\beta\hat{p}(v)/B}$. At $C = 0$, the uncapped offer is Δ/α. For the same marginal value, a larger C lowers that offer, limiting how much of the final payment budget this seller can claim at the expense of later sellers. Consequently, the rule balances the "blunted marginal surplus" $f(v \mid A_j) - \alpha\,\hat{p}(v)$ against the induced budget opportunity cost, discouraging excessive payments and favoring agents with high marginal value relative to their required compensation. Finally, since $\alpha > 1$, the blunted marginal surplus reserves a margin between value and payment, which ensures that the resulting social welfare remains at least a positive fraction of the value of the winner set.

5 Performance Analysis of Our Budget-Feasible Procurement Auction

In this section, we analyze the performance of our clock auction. We first establish, in Section 5.1, several general properties of the candidate winner sets maintained by Algorithm 1 that hold for all parameter settings. We then derive the approximation guarantees of Algorithm 1 for different choices of the parameters l and θ in Sections 5.2–5.4. To maintain the flow of the presentation, we defer most proofs and some auxiliary lemmas/corollaries to Appendix A and Appendix B.

5.1 General Properties of the Candidate Winner Sets

Recall that Algorithm 1 maintains l candidate winner sets $\{A_i : i \in [l]\}$ consisting of sellers who accept the offered prices. Therefore, the performance loss of Algorithm 1 largely stems from the exclusion of some sellers in the optimal solution S^* from these candidate winner sets. Quantifying this loss is thus useful for our analysis. According to Lines 4–9 of Algorithm 1, each seller $v \in S^*$ who is routed to some A_j ($j \in [l]$) with positive marginal value but is not included in A_j must belong to one of the following sets:

- R_j, the set of sellers in S^* who are routed to A_j but reject the offered price;
- D_j, the set of sellers in S^* who are routed to A_j, accept the offered price, but are discarded in Line 9 with probability $1 - \theta$. Clearly, $D_j = \emptyset$ when $\theta = 1$.

In what follows, we analyze the performance loss contributed by each of these parts.

Bounding the Loss from R_j: We first analyze the performance loss caused by the sets R_j, for all $j \in [l]$. Since each seller $v \in R_j$ rejects the offered price, the pricing rule in Algorithm 1 allows us to upper-bound the marginal value of v with respect to A_j, as formalized in Lemma 5.1:

Lemma 5.1. *Let $A_{i,v}$ be defined as in Observation 4.1, for each $i \in [l]$ and each $v \in N$. We have*

$$\forall j \in [l] : \sum\nolimits_{v \in R_j} f(v \mid A_{j,v}) \le \alpha\, c(R_j) + [f(A_j) - p(A_j)] \sum\nolimits_{v \in R_j} \left(e^{\beta c(v)/B} - 1\right) \tag{4}$$

Lemma 5.1 immediately yields an upper bound on the total marginal value of the sellers in $\cup_{j\in[l]} R_j$, which can be viewed as the performance loss caused by these sellers, as stated in Corollary 5.1. The corollary can be readily proved by using the inequality $\sum_{v\in S^*}(e^{\beta c(v)/B}-1) \le e^{\beta}-1$, which is established in Lemma A.1 of Appendix A.

Corollary 5.1. *We have* $\sum_{j\in[l]}\sum_{v\in R_j} f(v \mid A_{j,v}) \le \alpha\, c(S^*) + [f(A^*)-p(A^*)]\left(e^{\beta}-1\right)$.

Bounding the Loss from D_j: Next, we analyze the performance loss caused by the sellers in D_j, for any $j \in [l]$. Intuitively, the expected total loss caused by sellers in D_j can be upper-bounded by the expected value of A_j, because each seller v who accepts the price is included in the candidate winner set with probability θ, conditional on the random choices made before v is processed. The following lemma formalizes this intuition.

Lemma 5.2. *Let the random variable* $X_{j,v} \in \{0,1\}$ *for each* $v \in N$, *such that* $X_{j,v}=1$ *if* $v \in D_j$ *and* $X_{j,v}=0$ *otherwise. We have* $\mathbb{E}[\sum_{v\in N} X_{j,v}\cdot f(v \mid A_{j,v})] \le \frac{1-\theta}{\theta}\mathbb{E}[f(A_j)]$ *for all* $j \in [l]$.

With the above lemmas and corollaries, we are able to bound the performance loss caused by $S^* \setminus A^*$ with respect to A^*. However, since the final winner set S returned by Algorithm 1 may be a subset of A^*, it remains to relate S to A^* in order to derive the approximation ratio. To this end, we first use our pricing rule to show that the surplus of the candidate winner set grows exponentially with the payments, as formalized in the following lemma:

Lemma 5.3. *For any subset* $T=\{u_1,\ldots,u_m\} \subseteq A^*$ *of consecutively added sellers, we have*

$$f(T \mid Q) - p(T) \ge \big(f(Q)-p(Q)\big)\big(e^{\beta p(T)/B}-1\big), \tag{5}$$

where Q *denotes the set of sellers added to* A^* *before all sellers in* T.

Lemma 5.3 enables us to upper-bound the loss caused by removing certain sellers from A^* when constructing the final winner set S in Lines 16–17 of Algorithm 1. As a consequence, we can further upper-bound the surplus of A^* in terms of $f(S)-c(S)$, as stated in Lemma 5.4.

Lemma 5.4. *When Algorithm 1 finishes, we have*

$$f(A^*)-p(A^*) \le 2\left(1+\frac{1}{e^{\beta}-1}\right)[f(S)-c(S)] \tag{6}$$

5.2 Deterministic Approximation Guarantee for Monotone Submodular Valuations

In this section, we consider Algorithm 1 with parameters $l=1$ and $\theta=1$. Under this setting, the algorithm reduces to a deterministic clock auction for monotone submodular valuation functions $f(\cdot)$. Moreover, since $\theta=1$ implies $D_1=\emptyset$, the only performance loss arising from sellers in $S^* \setminus A^*$ is due to R_1. Building on the results in Section 5.1, we first establish the following lower bound on the surplus of the candidate winner set A^*:

Lemma 5.5. *When* $f(\cdot)$ *is a monotone submodular function, by setting* $l=1$ *and* $\theta=1$, *the set* $A^*=A_1$ *in Line 14 of Algorithm 1 satisfies* $\left[e^{\beta}+\frac{1}{\alpha-1}\right][f(A^*)-p(A^*)] \;\ge\; f(S^*)-\alpha\, c(S^*)$.

Lemma 5.4 and Lemma 5.5 together connect $f(S)-c(S)$ with $f(S^*)-\alpha c(S^*)$. Consequently, obtaining a strong approximation ratio for Algorithm 1 reduces to choosing suitable α and β to solve a nonlinear optimization problem, which can be addressed numerically, as shown in Theorem 5.1.

Theorem 5.1. *By setting* $l=1, \theta=1, \alpha \approx 8.523338$, *and* $\beta \approx 0.724829$, *Algorithm 1 achieves an* 8.5234*-approximation for monotone submodular valuation functions. That is, it outputs a set* S *satisfying* $f(S)-c(S) \ge \frac{1}{8.5234}f(S^*)-c(S^*)$, *using at most* $n+2$ *value oracle queries.*

Proof. Let $d=e^{\beta}+\frac{1}{\alpha-1}$ and $q=1+\frac{1}{e^{\beta}-1}$. Due to Lemma 5.4 and Lemma 5.5, we immediately get $2qd[f(S)-c(S)] \ge f(S^*)-\alpha\, c(S^*)$. Therefore, under the condition that $2qd \ge \alpha$, we achieves a $(2qd)$-approximation due to

$$f(S)-c(S) \ge f(S^*)/(2qd) - \alpha\, c(S^*)/(2qd) \ge f(S^*)/(2qd) - c(S^*) \tag{7}$$

Hence, optimizing the approximation ratio (i.e., $2qd$) reduces to solving the following problem:

$$\begin{aligned} \min_{\alpha,\beta} \quad & 2\left(e^{\beta}+(\alpha-1)^{-1}\right)\left(1+(e^{\beta}-1)^{-1}\right) \\ \text{s.t.} \quad & \alpha \le 2\left(e^{\beta}+(\alpha-1)^{-1}\right)\left(1+(e^{\beta}-1)^{-1}\right); \\ & \alpha > 1, \beta > 0; \end{aligned}$$

It can be verified that the optimal value of the above problem is at most 8.5234 by setting $\alpha \approx 8.523338$ and $\beta \approx 0.724829$. Note that each seller incurs at most 1 value oracle query due to Line 4 and Line 6, and Line 17 incurs 2 value oracle queries if it is executed. So Algorithm 1 incurs at most $n+2$ value oracle queries under the setting of $l=1$. This completes the proof. □

5.3 Expected Approximation Guarantee for Non-Monotone Submodular Valuations

In this section, we consider the case that inputting $l=1$ and $\theta \in (0,1)$ to Algorithm 1. In this case, Algorithm 1 becomes a randomized clock auction to handle a non-monotone submodular valuation function $f(\cdot)$. To analyze the approximation ratio under this case, we need to consider the performance loss cause by both R_1 and D_1 defined in Section 5.1, which are available through Corollary 5.1 and Lemma 5.2. So we can get Lemma 5.6:

Lemma 5.6. *When $f(\cdot)$ is a non-monotone submodular function, by setting $l=1$ and $\theta \in (0,1)$, the set $A^*=A_1$ in Line 14 of Algorithm 1 satisfies*

$$\left[e^{\beta}-1+\frac{\alpha}{\theta(\alpha-1)}\right]\mathbb{E}[f(A^*)-p(A^*)] \;\ge\; (1-\theta)f(S^*)-\alpha\, c(S^*) \tag{8}$$

The proof of Lemma 5.6 is broadly similar in spirit to that of Lemma 5.5. However, it additionally relies on the bound $\mathbb{E}[f(S^* \cup A_1)] \ge (1-\theta)f(S^*)$, established in Corollary A.1 in Appendix A. This additional step is necessary because, when $f(\cdot)$ is non-monotone, $f(S^* \cup A_1)$ cannot be used as an upper bound on $f(S^*)$, unlike in the proof of Lemma 5.5.

Combining Lemma 5.4 and Lemma 5.6, we derive the expected approximation ratio of Algorithm 1, as stated in Theorem 5.2. As in the proof of Theorem 5.1, establishing Theorem 5.2 also involves solving an optimization problem to choose suitable parameters θ, α, and β.

Theorem 5.2. *By setting $l=1, \theta \approx 0.3752, \alpha \approx 14.494$ and $\beta \approx 0.9907$, Algorithm 1 achieves an expected approximation ratio of 23.2 for non-monotone submodular valuations, That is, it outputs a set S satisfying $\mathbb{E}[f(S)-c(S)] \ge \frac{1}{23.2}f(S^*)-c(S^*)$, using at most $n+2$ value oracle queries.*

5.4 Deterministic Approximation Guarantee for Non-Monotone Submodular Valuations

In this section, we consider Algorithm 1 with parameters $l=2$ and $\theta=1$. Under this setting, the algorithm becomes a deterministic clock auction for a non-monotone submodular valuation function $f(\cdot)$. Analyzing the approximation ratio in this case is more involved than in Section 5.2, because we must account for the performance loss caused by $S^* \setminus A_j$ for each $j \in [l]$. Moreover, $S^* \setminus A_j$ contains not only the sellers in $R_1 \cup R_2$ (note that $D_1 \cup D_2 = \emptyset$ in this case), but also sellers in $S^* \cap A_i$, where $i \in [l] \setminus \{j\}$. Fortunately, the performance loss caused by $S^* \cap A_i$ can be upper-bounded by $f(A_i)$ due to the greedy routing rule in Line 4 of Algorithm 1. Using this idea, we obtain the following lemma through a careful analysis.

Lemma 5.7. *Suppose that $f(\cdot)$ is a general submodular function. By setting $l=2$ and $\theta=1$, the set A^* in Line 14 of Algorithm 1 satisfies $2\left[e^{\beta}+\frac{\alpha+1}{\alpha-1}\right][f(A^*)-p(A^*)] \;\ge\; f(S^*)-2\alpha\, c(S^*)$.*

With Lemma 5.7, we can use similar reasoning as that in Theorem 5.1 to get:

Theorem 5.3. *By setting $l=2, \theta=1, \alpha \approx 12.258$ and $\beta \approx 0.9066$, Algorithm 1 can achieve an approximation ratio of 24.52 for a non-monotone submodular valuation function. That is, it outputs a set S satisfying $f(S)-c(S) \ge \frac{1}{24.52}f(S^*)-c(S^*)$, using at most $2n+2$ value oracle queries.*

6 Performance Evaluation

In this section, we experimentally evaluate the effectiveness and efficiency of our clock auctions. All experiments were conducted on a CPU-only Apple M4 Pro workstation with 24 GB of memory.

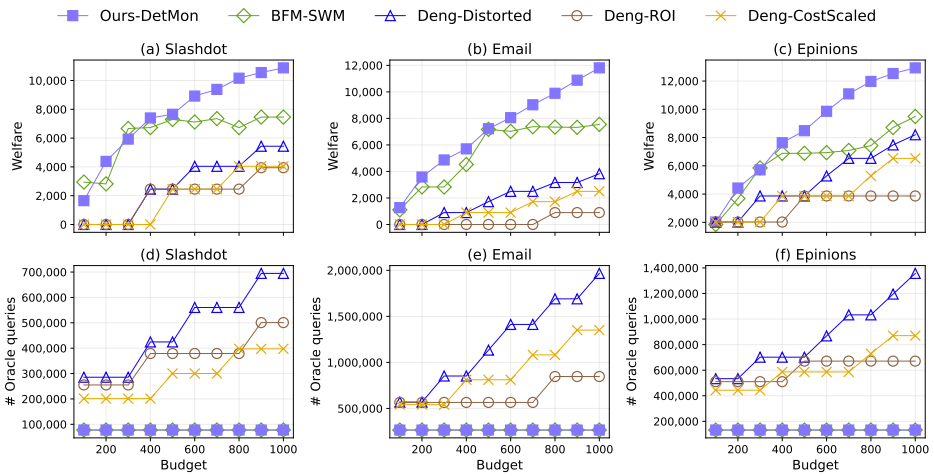

Figure 1: Comparison of the implemented mechanisms in terms of social welfare and the number of value-oracle queries for the influence-maximization application.

Since Cui et al. [13] is the only existing work that studies the same problem setting, we follow their application scenarios and experimental settings for a fair comparison.

For monotone submodular valuations, we follow Cui et al. [13] to consider an influence-maximization application on a directed social network $G = (V, E)$, where V and E denote the sets of nodes and edges, respectively. A buyer with budget B seeks to select a seed set $S \subseteq V$. Each node $v \in V$ has a private cost $c(v)$ and requires a payment of at least $c(v)$ from the buyer to serve as a seed node. The objective is to maximize the social welfare $f(S) - c(S)$, where $f(S)$ denotes the influence spread of S. Following [12, 13, 24, 28], we adopt the one-hop influence function $f(S) = |\{v \in V : \exists\, u \in S \text{ such that } (u, v) \in E\}|$ (this function is known to be monotone and submodular), and set the cost of each node $v \in V$ to $1 + \sqrt{d(v)}$, where $d(v)$ denotes the out-degree of v. As in Cui et al. [13], we use three public social-network datasets from SNAP [30]: SLASHDOT (77,360 nodes, 828,161 edges), EMAIL-EUALL (265,214 nodes, 418,956 edges), and EPINIONS (131,828 nodes, 840,799 edges).

We implement the following algorithms for comparison in the influence-maximization application:

- Our Algorithm 1, denoted by **OURS-DETMON**, with the input parameters specified in Theorem 5.1;
- The **BFM-SWM** mechanism of Cui et al. [13], which is the only existing auction mechanism for our problem. We adopt the best-performing parameter setting reported in Cui et al. [13] for monotone submodular valuations (cf. Section 6.1 therein);
- The mechanism of Deng et al. [15], originally proposed for monotone submodular welfare maximization under a cardinality constraint or without constraints. Since this mechanism does not account for budget constraints on payments, we follow the adaptation of Cui et al. [13] as follows. Specifically, we instantiate the framework of Deng et al. [15] with several regularized submodular maximization algorithms, namely Distorted Greedy [24], ROI Greedy [27], and Cost-Scaled Greedy [39], yielding three variants denoted by **DENG-DISTORTED**, **DENG-ROI**, and **DENG-COSTSCALED**, respectively. For each variant, we return the longest prefix of the selected solution that satisfies the budget constraint. Finally, we adopt the acceleration strategy and parameter settings reported in Deng et al. [15], following the implementation protocol of Cui et al. [13].

Experimental Results. Figure 1 reports two main metrics, namely social welfare and the number of value-oracle queries, for all implemented mechanisms on the three social-network datasets under varying budget levels. We observe that DENG-DISTORTED, DENG-ROI , and DENG-COSTSCALED perform the worst on both metrics, which is unsurprising since they are not specifically designed for budget-feasible social welfare maximization. In contrast, OURS-DETMON achieves higher social welfare than BFM-SWM in most cases, with improvements of up to 72.1%. Across all 30 dataset–budget pairs in Figure 1, OURS-DETMON achieves an average social welfare of 7864, representing a 29.3% improvement over the average social welfare of 6082 achieved by BFM-SWM. This result is

consistent with the fact that OURS-DETMON has a better approximation ratio than BFM-SWM. We also observe that BFM-SWM incurs a similar, albeit slightly larger, number of value-oracle queries than OURS-DETMON. This suggests that BFM-SWM may have comparable practical efficiency in this application, even though its theoretical complexity bound indicates that it can perform worse in other settings. Overall, the results in Figure 1 show that OURS-DETMON achieves the best social welfare while incurring the fewest value-oracle queries.

Additional Experiments. We also evaluate our approach on a crowdsourcing application with non-monotone submodular valuations. The results show that Algorithm 1 significantly outperforms the mechanism of Cui et al. [13] in terms of both social welfare and query complexity. Due to space limitations, details of the experimental setup and additional results for the crowdsourcing application are deferred to Appendix C.

7 Conclusion and Future Work

We studied the problem of budget-feasible social-welfare maximization with submodular valuations. For both monotone and non-monotone submodular valuations, we proposed simple yet effective single-round clock auctions that achieve improved approximation guarantees, lower value-oracle query complexity, and fewer pricing rounds than prior work. Experimental results on influence maximization and crowdsourcing further support the practical effectiveness and efficiency of our approach. Although our results improve upon prior approximation guarantees, an interesting open problem, which is also a limitation of our work, is whether these guarantees are tight.

For budget-feasible maximization of the buyer's value, Chen et al. [10] proved lower bounds of $1 + \sqrt{2}$ for deterministic truthful mechanisms and 2 for universally truthful randomized mechanisms, even with additive valuations. Their objective is the public valuation $f(S)$. In our model, $f(S) - c(S)$ depends on private costs and can be negative, and the benchmark deducts the full cost of the optimal set. The bounds for value maximization therefore do not directly establish a lower bound on our problem. We plan to investigate the lower bounds for our problem in future work.

Acknowledgments and Disclosure of Funding

This work was supported in part by the Special Project for Follow-up Development Research at Shanghai University of Finance and Economics (Project No. 2026110608), in part by the Fundamental Research Funds for the Central Universities (Project No. 2025110402), in part by National Key R&D Program of China under Grant No. 2024YFA1012700, in part by the National Natural Science Foundation of China (NSFC) under Grant No. 62402410, in part by Guangdong Provincial Project (No. 2023QN10X025), in part by NSFC under Grant 62472253, and in part by the Taishan Scholars Program under Grant tsqn202408025.

References

[1] Georgios Amanatidis, Pieter Kleer, and Guido Schäfer. Budget-feasible mechanism design for non-monotone submodular objectives: Offline and online. *Mathematics of Operations Research*, 47(3): 2286–2309, 2022.

[2] Nima Anari, Gagan Goel, and Afshin Nikzad. Mechanism design for crowdsourcing: An optimal 1-1/e competitive budget-feasible mechanism for large markets. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 266–275, 2014.

[3] Lawrence M. Ausubel, Peter Cramton, and Paul R. Milgrom. The clock-proxy auction: A practical combinatorial auction design. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, pages 115–138. MIT Press, Cambridge, MA, 2006.

[4] Eric Balkanski, Pranav Garimidi, Vasilis Gkatzelis, Daniel Schoepflin, and Xizhi Tan. Deterministic budget-feasible clock auctions. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2940–2963, 2022.

[5] Eric Balkanski, Pranav Garimidi, Vasilis Gkatzelis, Daniel Schoepflin, and Xizhi Tan. Deterministic budget-feasible clock auctions. *Operations Research*, 2025.

[6] 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.

[7] Jeff Bilmes. Submodularity in machine learning and artificial intelligence. *arXiv:2202.00132*, 2022.

[8] Sandro Brusco and Giuseppe Lopomo. Collusion via signalling in simultaneous ascending bid auctions with heterogeneous objects, with and without complementarities. *Review of Economic Studies*, 69(2): 407–436, 2002.

[9] Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. Submodular maximization with cardinality constraints. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1433–1452, 2014.

[10] Ning Chen, Nick Gravin, and Pinyan Lu. On the approximability of budget feasible mechanisms. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 685–699, 2011.

[11] Peter Cramton and Jesse A. Schwartz. Collusive bidding: Lessons from the FCC spectrum auctions. *Journal of Regulatory Economics*, 17(3):229–252, 2000.

[12] Shuang Cui, Kai Han, Jing Tang, and He Huang. Constrained subset selection from data streams for profit maximization. In *Proceedings of the 32nd ACM Web Conference (WWW)*, pages 1822–1831, 2023.

[13] Shuang Cui, He Huang, Yu-e Sun, and Chen Xue. Budget-feasible mechanisms for submodular welfare maximization in procurement auctions. In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, 2026. arXiv:2605.00411.

[14] Pranav Dandekar, Nadia Fawaz, and Stratis Ioannidis. Privacy auctions for recommender systems. *ACM Transactions on Economics and Computation*, 2(3):12:1–12:22, 2014.

[15] Yuan Deng, Amin Karbasi, Vahab Mirrokni, Renato Paes Leme, Grigoris Velegkas, and Song Zuo. Procurement auctions via approximately optimal submodular optimization. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, volume 267, pages 13152–13180, 2025.

[16] Paul Duetting, Federico Fusco, Silvio Lattanzi, Ashkan Norouzi-Fard, and Morteza Zadimoghaddam. Deletion robust submodular maximization over matroids. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, volume 162, pages 5671–5693, 2022.

[17] Marwa El Halabi, Federico Fusco, Ashkan Norouzi-Fard, Jakab Tardos, and Jakub Tarnawski. Fairness in streaming submodular maximization over a matroid constraint. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, volume 202, pages 9150–9171, 2023.

[18] Matthew Fahrbach, Vahab Mirrokni, and Morteza Zadimoghaddam. Non-monotone submodular maximization with nearly optimal adaptivity and query complexity. In *International Conference on Machine Learning (ICML)*, pages 1833–1842, 2019.

[19] Moran Feldman, Christopher Harshaw, and Amin Karbasi. Greed is good: Near-optimal submodular maximization via greedy optimization. In *Conference on Learning Theory (COLT)*, pages 758–784, 2017.

[20] Nick Gravin, Yaonan Jin, Pinyan Lu, and Chenhao Zhang. Optimal budget-feasible mechanisms for additive valuations. *ACM Transactions on Economics and Computation*, 8(4):21:1–21:15, 2020.

[21] Kai Han and Pinyan Lu. Improved budget-feasible mechanisms for submodular valuations: Beating 4 deterministically in linear time. In *International Conference on Web and Internet Economics (WINE)*, 2025.

[22] Kai Han, You Wu, He Huang, and Shuang Cui. Triple eagle: Simple, fast and practical budget-feasible mechanisms. In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, volume 36, pages 33894–33911, 2023.

[23] Kai Han, Haotian Zhang, and Shuang Cui. Efficient and effective budget-feasible mechanisms for submodular valuations. *Artificial Intelligence*, 345:104348, 2025.

[24] Christopher Harshaw, Moran Feldman, Justin Ward, and Amin Karbasi. Submodular maximization beyond non-negativity: Guarantees, fast algorithms, and applications. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97, pages 2634–2643, 2019.

[25] He Huang, Kai Han, Shuang Cui, and Jing Tang. Randomized pricing with deferred acceptance for revenue maximization with submodular objectives. In *Proceedings of the 2023 ACM Web Conference (WWW)*, pages 3530–3540, 2023.

[26] Pooya Jalaly and Éva Tardos. Simple and efficient budget feasible mechanisms for monotone submodular valuations. *ACM Transactions on Economics and Computation (TEAC)*, 9(1):1–20, 2021.

[27] Tianyuan Jin, Yu Yang, Renchi Yang, Jieming Shi, Keke Huang, and Xiaokui Xiao. Unconstrained submodular maximization with modular costs: Tight approximation and application to profit maximization. *Proceedings of the VLDB Endowment (PVLDB)*, 14(10):1756–1768, 2021.

[28] Ehsan Kazemi, Shervin Minaee, Moran Feldman, and Amin Karbasi. Regularized submodular maximization at scale. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139, pages 5356–5366, 2021.

[29] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Technical report, Citeseer*, 2009.

[30] Jure Leskovec and Andrej Krevl. SNAP datasets: Stanford large network dataset collection. `https://snap.stanford.edu/data`, 2014.

[31] Jonathan Levin and Andrzej Skrzypacz. Properties of the combinatorial clock auction. *American Economic Review*, 106(9):2528–2551, 2016.

[32] Xiang Liu, Hau Chan, Minming Li, Weiwei Wu, and Yingchao Zhao. Budget-feasible mechanisms for proportionally selecting agents from groups. *Artificial Intelligence*, 323:103975, 2023.

[33] Yuan Luo and Nicholas R. Jennings. A budget-limited mechanism for category-aware crowdsourcing of multiple-choice tasks. *Artificial Intelligence*, 299:103538, 2021.

[34] Robert C. Marshall and Leslie M. Marx. The vulnerability of auctions to bidder collusion. *Quarterly Journal of Economics*, 124(2):883–910, 2009.

[35] Paul Milgrom and Ilya Segal. Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, 128(1):1 – 31, 2020.

[36] Baharan Mirzasoleiman, Ashwinkumar Badanidiyuru, and Amin Karbasi. Fast constrained submodular maximization: Personalized data summarization. In *International Conference on Machine Learning (ICML)*, pages 1358–1367, 2016.

[37] Rian Neogi, Kanstantsin Pashkovich, and Chaitanya Swamy. Budget-feasible mechanism design: Simpler, better mechanisms and general payment constraints. In *Proceedings of the 15th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 287, pages 84:1–84:22, 2024.

[38] Rian Neogi, Kanstantsin Pashkovich, and Chaitanya Swamy. An O(log log n)-approximate budget feasible mechanism for subadditive valuations. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC)*, 2025.

[39] Sofia Maria Nikolakaki, Alina Ene, and Evimaria Terzi. An efficient framework for balancing submodularity and cost. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1256–1266, 2021.

[40] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007. ISBN 9780511800481.

[41] Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.

[42] Yaron Singer. Budget feasible mechanisms. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 765–774, 2010.

[43] Yaron Singer. How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In *Proceedings of the 5th ACM International Conference on Web Search and Data Mining (WSDM)*, pages 733–742, 2012.

[44] Yaron Singer and Manas Mittal. Pricing mechanisms for crowdsourcing markets. In *Proceedings of the 22nd International Conference on World Wide Web (WWW)*, pages 1157–1166, 2013.

[45] Zhenyu Wen, Wanglei Feng, Di Wu, Haozhen Hu, Chang Xu, Bin Qian, Zhen Hong, Cong Wang, and Shouling Ji. Flmarket: Enabling privacy-preserved pre-training data pricing for federated learning. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1587–1598, 2025.

[46] Jingwen Zhang, Yuezhou Wu, and Rong Pan. Incentive mechanism for horizontal federated learning based on reputation and reverse auction. In *Proceedings of the 2021 ACM Web Conference (WWW)*, pages 947–956, 2021.

[47] Zhenzhe Zheng, Fan Wu, Xiaofeng Gao, Hongzi Zhu, Shaojie Tang, and Guihai Chen. A budget feasible incentive mechanism for weighted coverage maximization in mobile crowdsensing. *IEEE Transactions on Mobile Computing*, 16(9):2392–2407, 2017.

A Auxiliary Lemmas and Corollaries

Lemma A.1. *For any $\beta > 0$, we have $\sum_{v \in S^*} \left(e^{\beta c(v)/B} - 1\right) \leq e^{\beta} - 1$.*

Proof. For any $v \in S^*$, let $y(v) = e^{\beta c(v)/B} - 1$. Since $\beta \geq 0$ and $c(v) \geq 0$, we have $y(v) \geq 0$. Moreover, as $c(S^*) \leq B$, we get

$$e^{\beta} \geq e^{\beta c(S^*)/B} = e^{\beta \sum_{v \in S^*} c(v)/B} = \prod_{v \in S^*} e^{\beta c(v)/B} = \prod_{v \in S^*} (1 + y(v)) \geq 1 + \sum_{v \in S^*} y(v)$$

which proves the lemma. □

Lemma A.2. *The pricing rule in Line 6 of Algorithm 1, i.e., $f(v \mid A_j) - \alpha \cdot \hat{p}(v) = (f(A_j) - p(A_j)) \left[e^{\beta \hat{p}(v)/B} - 1\right]$, admits a unique solution $\hat{p}(v) \in \left(0, f(v \mid A_j)/\alpha\right]$,*

Proof. Let $x = \hat{p}(v)$, and let

$$\phi(x) = f(v \mid A_j) - \alpha x - \big(f(A_j) - p(A_j)\big) \left(e^{\beta x/B} - 1\right).$$

Line 5 of Algorithm 1 ensures $\phi(0) = f(v \mid A_j) > 0$. Moreover, by Observation 4.1, we have $f(A_j) - p(A_j) \geq 0$ and hence

$$\phi\left(f(v \mid A_j)/\alpha\right) = -\big(f(A_j) - p(A_j)\big) \left(e^{\beta f(v|A_j)/(\alpha B)} - 1\right) \leq 0.$$

Furthermore,

$$\phi'(x) = -\alpha - \big(f(A_j) - p(A_j)\big)(\beta/B)e^{\beta x/B} < 0$$

for all $x \geq 0$. Hence ϕ is strictly decreasing on $\mathbb{R}_{\geq 0}$, and therefore has a unique root in $\big(0, f(v \mid A_j)/\alpha\big]$. This root is $\hat{p}(v)$. □

Lemma A.3 (cf. Lemma 2.2 of [9])**.** *Let N be a ground set, and let $g : 2^N \to \mathbb{R}_{\geq 0}$ be any non-negative submodular function. Suppose Y is a random subset of N such that each element $e \in N$ is included in Y with probability at most θ, not necessarily independently. Then,*

$$\mathbb{E}[g(Y)] \geq (1 - \theta)g(\emptyset).$$

Corollary A.1. *Let A_1 be the candidate winner set constructed in Algorithm 1 when $l = 1$ and $\theta \in (0, 1)$, we have $\mathbb{E}[f(S^* \cup A_1)] \geq (1 - \theta)f(S^*)$.*

Proof. This corollary follows directly from Lemma 2.2 of [9] stated above. We include the proof for completeness. Define $g : 2^N \to \mathbb{R}_{\geq 0}$ by $g(X) = f(X \cup S^*)$. Since f is non-negative and submodular, so is g. Each seller enters A_1 only if it accepts its price offer and passes the independent retention step of probability θ; hence every element is included in A_1 with marginal probability at most θ. Applying Lemma A.3 gives

$$\mathbb{E}[f(S^* \cup A_1)] = \mathbb{E}[g(A_1)] \geq (1 - \theta)g(\emptyset) = (1 - \theta)f(S^*),$$

which completes the proof. □

B Missing Proofs

B.1 Proof of Lemma 4.1

Proof. (1) We first prove budget-feasibility and NAS. In the case of $p(A^*) \leq B$, we have $S = A^*$ due to Line 14 of Algorithm 1. Therefore, we get $p(S) \leq B$ and

$$f(S) - c(S) \geq f(S) - p(S) \geq 0,$$

where the first inequality is due to $c(v) \leq p(v)$ for any $v \in A^*$, and the second inequality is due to Observation 4.1. On the other hand, if $p(A^*) > B$, we have either $S = \{w\}$ or $S = H \setminus \{w\}$ due to

Lines 16-17 of Algorithm 1. Note that $p(w) \le B$ due to $p(w) = \min\{B, \hat{p}(w)\}$ and $p(H \setminus \{w\}) \le B$ due to the definition of H. So $p(S) \le B$. Moreover, suppose that $A^* = A_j$, we can get

$$f(w) - p(w) \ge f(w \mid A_{j,w}) - p(w) \ge 0$$
$$f(H \setminus \{w\}) - p(H \setminus \{w\}) \ge f(H \setminus \{w\} \mid \{w\} \cup A_{j,w}) - p(H \setminus \{w\}) \ge 0$$

due to the submodularity of $f(\cdot)$ and Observation 4.1. So we always have $f(S) - c(S) \ge f(S) - p(S) \ge 0$, which completes the proof on budget-feasibility and NAS.

(2) Next, we prove truthfulness and IR. Since Algorithm 1 is a one-round clock auction (see Section 3 for the definition of a clock auction), it naturally satisfies truthfulness and individual rationality [4, 35]. We provide the proof below for completeness and to assist readers who may be unfamiliar with clock auctions.

Fix a seller v, the activity strategies of all other sellers, and the realization of the mechanism's internal randomness. When v is processed, the price $p(v)$, if an offer is made to v, is computed from the current candidate winner sets, the valuation oracle, the budget, and the mechanism parameters before v's accept/reject decision. Hence $p(v)$ is independent of v's current decision and of v's private cost. If no price offer is made to v, then v cannot be selected as a winner and obtains utility 0.

Suppose an price offer $p(v)$ is made. If v rejects, it permanently exits and obtains utility 0. If v accepts, then, regardless of how this acceptance affects the subsequent candidate winner sets and the final winner selection rule, v's utility is either 0 if it is not selected as a winner, or $p(v) - c(v)$ if it is selected. Therefore, if $c(v) \le p(v)$, accepting weakly dominates rejecting; if $c(v) > p(v)$, rejecting weakly dominates accepting. Thus the dominant activity strategy is to accept if and only if $c(v) \le p(v)$, which establishes truthfulness of Algorithm 1.

Under this truthful activity strategy, every selected seller receives a payment at least its true cost and therefore obtains non-negative utility, while every non-selected seller obtains utility 0. This proves individual rationality. So the lemma follows. □

B.2 Proof of Lemma 5.1

Proof. Since $R_j \subseteq S^*$, every $v \in R_j$ satisfies $c(v) \le B$. Because v rejects the offer, we have $c(v) > p(v) = \min\{B, \hat{p}(v)\}$. This implies $\hat{p}(v) < B$ and hence $p(v) = \hat{p}(v) < c(v)$. Therefore, due to the pricing rule in Line 6 of Algorithm 1, for each $v \in R_j$, we have

$$\begin{aligned} & f(v \mid A_{j,v}) - \alpha\, c(v) < f(v \mid A_{j,v}) - \alpha\, \hat{p}(v) \\ = \;& [f(A_{j,v}) - p(A_{j,v})]\Big(e^{\beta \hat{p}(v)/B} - 1\Big) \le [f(A_j) - p(A_j)]\Big(e^{\beta c(v)/B} - 1\Big), \end{aligned} \tag{9}$$

where the second inequality follows from Observation 4.1. The lemma then follows by summing the above inequality for all $v \in R_j$ and re-arranging. □

B.3 Proof of Lemma 5.2

Proof. The proof is inspired in spirit by [19]; we include it for completeness. Let $\mathcal{E}_v$ denote the realization of the random choices made by Algorithm 1 immediately before seller $v \in N$ is processed, and let $Y_{j,v} \in \{0, 1\}$ be the indicator variable for the event $v \in A_j$, i.e., $Y_{j,v} = 1$ if $v \in A_j$ and $Y_{j,v} = 0$ otherwise. For any $v \in N$, conditioned on $\mathcal{E}_v$, if v is routed to A_j and accepts the offered price, then

$$\mathbb{E}[X_{j,v} \mid \mathcal{E}_v] = 1 - \theta \qquad \text{and} \qquad \mathbb{E}[Y_{j,v} \mid \mathcal{E}_v] = \theta; \tag{10}$$

otherwise we have $\mathbb{E}[X_{j,v} \mid \mathcal{E}_v] = \mathbb{E}[Y_{j,v} \mid \mathcal{E}_v] = 0$. Therefore, in all cases,

$$\mathbb{E}[X_{j,v} f(v \mid A_{j,v}) \mid \mathcal{E}_v] \le \frac{1-\theta}{\theta}\, \mathbb{E}[Y_{j,v} f(v \mid A_{j,v}) \mid \mathcal{E}_v]. \tag{11}$$

Summing over all $v \in N$ and taking total expectation, we obtain

$$\mathbb{E}\left[\sum_{v \in N} X_{j,v} f(v \mid A_{j,v})\right] \le \frac{1-\theta}{\theta}\, \mathbb{E}\left[\sum_{v \in N} Y_{j,v} f(v \mid A_{j,v})\right]. \tag{12}$$

Since $\sum_{v \in N} Y_{j,v} f(v \mid A_{j,v}) \le f(A_j)$, the claimed bound follows. □

B.4 Proof of Lemma 5.3

Proof. Let $T_i = \{u_1, \ldots, u_i\}$ for $i = 0, \ldots, m$, where $T_0 = \emptyset$. Denote $Q_i = Q \cup T_i$. Then

$$\begin{aligned}\forall i \in [m] : f(u_i \mid Q_{i-1}) - p(u_i) &\geq f(u_i \mid Q_{i-1}) - \alpha\hat{p}(u_i) \\ &= \big(f(Q_{i-1}) - p(Q_{i-1})\big)\left(e^{\beta\hat{p}(u_i)/B} - 1\right) \\ &\geq \big(f(Q_{i-1}) - p(Q_{i-1})\big)\left(e^{\beta p(u_i)/B} - 1\right), \end{aligned} \tag{13}$$

where the first inequality follows from $p(u_i) \leq \hat{p}(u_i)$ and $\alpha > 1$, and the equality follows from the pricing rule in Line 6 of Algorithm 1. Moreover, we have

$$\forall i \in [m] : f(Q_i) - p(Q_i) = f(Q_{i-1}) - p(Q_{i-1}) + f(u_i \mid Q_{i-1}) - p(u_i).$$

Hence, using Eqn. (13), we get

$$\forall i \in [m] : f(Q_i) - p(Q_i) \geq \big(f(Q_{i-1}) - p(Q_{i-1})\big)e^{\beta p(u_i)/B}.$$

Multiplying these inequalities for $i = 1, \ldots, m$ gives

$$f(Q_m) - p(Q_m) \geq \big(f(Q) - p(Q)\big)e^{\beta p(T)/B}.$$

Since

$$f(T \mid Q) - p(T) = f(Q_m) - p(Q_m) - \big(f(Q) - p(Q)\big),$$

we obtain

$$f(T \mid Q) - p(T) \geq \big(f(Q) - p(Q)\big)\left(e^{\beta p(T)/B} - 1\right).$$

This completes the proof. □

B.5 Proof of Lemma 5.4

Proof. Consider the set H in Line 16 of Algorithm 1. Using Lemma 5.3 and $p(H) > B$, we have

$$\begin{aligned} & f(H \mid A^* \setminus H) - p(H) \geq [f(A^* \setminus H) - p(A^* \setminus H)][e^{\beta p(H)/B} - 1] \\ \geq \quad & [f(A^* \setminus H) - p(A^* \setminus H)][e^{\beta} - 1] \end{aligned} \tag{14}$$

Using the above equality and the submodularity of $f(\cdot)$, we get

$$\begin{aligned} f(A^*) - p(A^*) &= [f(H \mid A^* \setminus H) - p(H)] + [f(A^* \setminus H) - p(A^* \setminus H)] \\ &\leq \left(1 + \frac{1}{e^{\beta} - 1}\right)[f(H \mid A^* \setminus H) - p(H)] \\ &\leq \left(1 + \frac{1}{e^{\beta} - 1}\right)[f(H) - p(H)]. \end{aligned} \tag{15}$$

Due to Lines 16-17 of Algorithm 1 and the submodularity of $f(\cdot)$, when $p(A^*) > B$, we have

$$\begin{aligned} f(H) - p(H) &\leq [f(H \setminus \{w\}) - p(H \setminus \{w\})] + [f(w) - p(w)] \\ &\leq 2[f(S) - p(S)] \leq 2[f(S) - c(S)] \end{aligned} \tag{16}$$

So the lemma follows by combining Eqn. (15) and Eqn. (16). On the other hand, when $p(A^*) \leq B$, we have $S = A^*$ due to Line 14 of Algorithm 1, so we have

$$\begin{aligned} f(A^*) - p(A^*) \quad &\leq \quad 2\left(1 + \frac{1}{e^{\beta} - 1}\right)[f(A^*) - p(A^*)] \\ &\leq \quad 2\left(1 + \frac{1}{e^{\beta} - 1}\right)[f(A^*) - c(A^*)] = 2\left(1 + \frac{1}{e^{\beta} - 1}\right)[f(S) - c(S)] \end{aligned}$$

where the second inequality is due to $p(v) \geq c(v)$ for any $v \in A^*$. So the lemma still follows. This completes the proof. □

B.6 Proof of Lemma 5.5

Proof. In this case, we can use the submodularity of $f(\cdot)$ and Corollary 5.1 to get

$$\begin{aligned} f(S^* \cup A_1) &\leq f(A_1) + \sum\nolimits_{v \in S^* \setminus A_1} f(v \mid A_{1,v}) \\ &\leq f(A_1) + \alpha\, c(S^*) + [f(A^*) - p(A^*)]\left(e^\beta - 1\right) \\ &= \alpha\, c(S^*) + e^\beta f(A^*) - [e^\beta - 1]p(A^*) \end{aligned} \tag{17}$$

Since $f(A^*) - \alpha\, p(A^*) \geq 0$ due to Observation 4.1, we get

$$0 \leq \frac{1}{\alpha - 1} f(A^*) - \frac{\alpha}{\alpha - 1} p(A^*) \tag{18}$$

Adding Eqn. (17) and Eqn. (18) yields

$$\left[e^\beta + \frac{1}{\alpha - 1}\right][f(A^*) - p(A^*)] \;\geq\; f(S^* \cup A_1) - \alpha\, c(S^*) \tag{19}$$

Finally, since $f(\cdot)$ is monotone, we have $f(S^* \cup A_1) \geq f(S^*)$. So the lemma follows. □

B.7 Proof of Lemma 5.6

Proof. Similar to Eqn. (17) in the proof of Lemma 5.5, we can use the submodularity of $f(\cdot)$ and Corollary 5.1 to get

$$\begin{aligned} f(S^* \cup A_1) &\leq f(A_1) + \sum_{v \in S^* \setminus A_1} f(v \mid A_{1,v}) \\ &\leq f(A_1) + \sum_{v \in R_1} f(v \mid A_{1,v}) + \sum_{v \in D_1} f(v \mid A_{1,v}) \\ &\leq f(A_1) + \alpha\, c(S^*) + [f(A^*) - p(A^*)]\left(e^\beta - 1\right) + \sum_{v \in N} X_{1,v} \cdot f(v \mid A_{1,v}) \\ &= \alpha\, c(S^*) + e^\beta f(A^*) - [e^\beta - 1]p(A^*) + \sum_{v \in N} X_{1,v} \cdot f(v \mid A_{1,v}) \end{aligned} \tag{20}$$

where R_1, D_1 are defined in Section 5.1, and the random variable $X_{1,v} \in \{0, 1\}$ is defined in Lemma 5.2. Taking expectations on both sides of the above equation and applying Lemma 5.2 and Corollary A.1, we obtain

$$(1 - \theta) f(S^*) - \alpha\, c(S^*) \leq \left(e^\beta + \frac{1 - \theta}{\theta}\right) \mathbb{E}[f(A^*)] - [e^\beta - 1]\mathbb{E}[p(A^*)] \tag{21}$$

Since $f(A^*) - \alpha\, p(A^*) \geq 0$ due to Observation 4.1, we get

$$0 \leq \frac{1}{\theta(\alpha - 1)} \mathbb{E}[f(A^*)] - \frac{\alpha}{\theta(\alpha - 1)} \mathbb{E}[p(A^*)] \tag{22}$$

Adding Eqn. (21) and Eqn. (22) yields

$$\left[e^\beta - 1 + \frac{\alpha}{\theta(\alpha - 1)}\right] \mathbb{E}[f(A^*) - p(A^*)] \;\geq\; (1 - \theta) f(S^*) - \alpha\, c(S^*) \tag{23}$$

So the lemma follows. □

B.8 Proof of Theorem 5.2

Proof. Let $z = e^\beta - 1 + \frac{\alpha}{\theta(\alpha - 1)}$ and $q = 1 + \frac{1}{e^\beta - 1}$. By combining Lemma 5.4 and Lemma 5.6, we obtain the following equation:

$$2zq \cdot \mathbb{E}[f(S) - c(S)] \geq (1 - \theta)\, f(S^*) - \alpha\, c(S^*)$$

Re-arranging the above equation, we get

$$\mathbb{E}[f(S) - c(S)] \geq \frac{1-\theta}{2zq} f(S^*) - \frac{\alpha}{2zq} c(S^*) \geq \frac{1-\theta}{2zq} f(S^*) - c(S^*)$$

under the condition that $2zq \geq \alpha$. Therefore, optimizing the approximation ratio $2zq/(1-\theta)$ reduces to solving following optimization problem:

$$\begin{aligned} \min_{\alpha,\beta,\theta} \quad & 2\left[e^{\beta} - 1 + \frac{\alpha}{\theta(\alpha-1)}\right]\left(1 + \frac{1}{e^{\beta}-1}\right)(1-\theta)^{-1} \\ \text{s.t.} \quad & \alpha \leq 2\left[e^{\beta} - 1 + \frac{\alpha}{\theta(\alpha-1)}\right]\left(1 + \frac{1}{e^{\beta}-1}\right); \\ & \alpha > 1, \beta > 0, \theta \in (0,1); \end{aligned}$$

It can be verified that the optimal value of the above optimization problem is at most 23.2 by setting $\theta \approx 0.3752, \alpha \approx 14.494, \beta \approx 0.9907$. Finally, it is easy to see that the maximum number of value-oracle queries incurred in this case is the same as in Theorem 5.1. This completes the proof. □

B.9 Proof of Lemma 5.7

Proof. Any seller $v \in S^* \setminus A_1$ whose marginal value is positive when it is processed must belong either to $S^* \cap A_2$ (if it is routed to A_2 and accepts the price) or to $S^* \setminus (A_1 \cup A_2)$ (if it rejects the price). Recall that Section 5.1 defines the set $R_j (\forall j \in [l])$ containing the sellers who are routed to A_j but reject the price. Therefore, Line 4 of Algorithm 1 guarantees that

$$\forall v \in S^* \cap A_2 \cup R_2 : f(v \mid A_{2,v}) \geq f(v \mid A_{1,v}) \tag{24}$$

Note that any seller $v \in S^*$ who is never offered a price must satisfy $f(v \mid A_i) \leq f(v \mid A_{i,v}) \leq 0$ for all $i \in [l]$ due to Line 5 of Algorithm 1. Using Eqn. (24) and the submodularity of $f(\cdot)$, we get

$$\begin{aligned} f(S^* \cup A_1) &\leq f(A_1) + \sum_{v \in S^* \cap A_2} f(v \mid A_{1,v}) + \sum_{v \in S^* \setminus (A_1 \cup A_2)} f(v \mid A_1) \\ &\leq f(A_1) + \sum_{v \in S^* \cap A_2} f(v \mid A_{2,v}) + \sum_{v \in R_1} f(v \mid A_{1,v}) + \sum_{v \in R_2} f(v \mid A_{1,v}) \\ &\leq f(A_1) + f(A_2) + \sum_{v \in R_1} f(v \mid A_{1,v}) + \sum_{v \in R_2} f(v \mid A_{2,v}) \\ &= \sum_{i \in \{1,2\}} \left[f(A_i) + \sum_{v \in R_i} f(v \mid A_{i,v}) \right] \end{aligned} \tag{25}$$

Recall that Lemma 5.1 guarantees:

$$\forall i \in \{1,2\} : \sum_{v \in R_i} f(v \mid A_{i,v}) \leq \alpha \sum_{v \in R_i} c(v) + [f(A_i) - p(A_i)] \sum_{v \in R_i} \left(e^{\beta c(v)/B} - 1 \right); \tag{26}$$

Let $M_i = \sum_{v \in R_i} \left(e^{\beta c(v)/B} - 1 \right)$. Combining Eqn. (25) and Eqn. (26) gives us

$$\begin{aligned} f(S^* \cup A_1) &\leq \sum_{i \in \{1,2\}} \left[f(A_i) + \sum_{v \in R_i} f(v \mid A_{i,v}) \right] \\ &\leq \sum_{i \in \{1,2\}} \left[f(A_i) + \alpha \sum_{v \in R_i} c(v) + [f(A_i) - p(A_i)] M_i \right] \\ &\leq \alpha\, c(S^*) + \sum_{i \in \{1,2\}} [f(A_i)(M_i + 1) - p(A_i) M_i] \end{aligned} \tag{27}$$

Due to Observation 4.1, we have $f(A_i) \geq \alpha p(A_i)$ for each $i \in \{1,2\}$, which implies

$$0 \leq \sum_{i \in \{1,2\}} \left[\frac{1}{\alpha - 1} f(A_i) - \frac{\alpha}{\alpha - 1} p(A_i) \right] \tag{28}$$

Summing Eqn. (27) and Eqn. (28), we get

$$\begin{aligned} f(S^* \cup A_1) - \alpha c(S^*) &\leq \sum_{i \in \{1,2\}} \left(M_i + \frac{\alpha}{\alpha - 1} \right) [f(A_i) - p(A_i)] \\ &\leq [f(A^*) - p(A^*)] \left(\frac{2\alpha}{\alpha - 1} + \sum_{i \in \{1,2\}} M_i \right) \\ &\leq [f(A^*) - p(A^*)] \left(e^\beta - 1 + \frac{2\alpha}{\alpha - 1} \right) \end{aligned} \tag{29}$$

Where the last inequality used Lemma A.1. By similar reasoning to the above, we can also get

$$f(S^* \cup A_2) - \alpha c(S^*) \leq [f(A^*) - p(A^*)] \left(e^\beta - 1 + \frac{2\alpha}{\alpha - 1} \right) \tag{30}$$

Since A_1 and A_2 are disjoint, the submodularity and non-negativity of $f(\cdot)$ ensures

$$f(S^* \cup A_1) + f(S^* \cup A_2) \geq f(S^* \cup A_1 \cup A_2) + f(S^*) \geq f(S^*) \tag{31}$$

The lemma then follows by combining Eqn. (29), Eqn. (30), and Eqn. (31). □

B.10 Proof of Theorem 5.3

Proof. Combining Lemma 5.4 and Lemma 5.7 gives us

$$4 \left[e^\beta + \frac{\alpha + 1}{\alpha - 1} \right] \left(1 + \frac{1}{e^\beta - 1} \right) [f(S) - c(S)] \;\geq\; f(S^*) - 2\alpha\, c(S^*) \tag{32}$$

which implies a $(4sq)$-approximation under the condition that $\alpha \leq 2sq$, with $s = e^\beta + \frac{\alpha+1}{\alpha-1}$ and $q = 1 + \frac{1}{e^\beta - 1}$. Optimizing this approximation ratio reduces to solving the following problem:

$$\begin{aligned} \min_{\alpha, \beta} \quad & 4 \left[e^\beta + \frac{\alpha + 1}{\alpha - 1} \right] \left(1 + \frac{1}{e^\beta - 1} \right) \\ \text{s.t.} \quad & \alpha \leq 2 \left[e^\beta + \frac{\alpha + 1}{\alpha - 1} \right] \left(1 + \frac{1}{e^\beta - 1} \right); \\ & \alpha > 1, \beta > 0; \end{aligned}$$

It can be verified that the optimal value of the above optimization problem is at most 24.5161 by setting $\alpha \approx 12.258$ and $\beta \approx 0.9066$. Note that each seller incurs 2 value oracle queries due to Line 4 and Line 6, and Line 17 incurs 2 value oracle queries if it is executed. So Algorithm 1 incurs at most $2n + 2$ value oracle queries under the setting of $l = 2$. This completes the proof. □

C Additional Experiments

Following Cui et al. [13] and [23, 25, 26], we consider a crowdsourcing application in which a buyer with budget B procures images from a set N of workers (i.e., sellers), each of whom owns one image and incurs a private cost. For a selected set $S \subseteq N$, its overall quality is defined as

$$f(S) = \frac{1}{|N|} \left[\sum_{u \in N} \sum_{v \in S} \mathrm{Sim}(u, v) - \sum_{u \in S} \sum_{v \in S} \mathrm{Sim}(u, v) \right], \tag{33}$$

where $\mathrm{Sim}(u, v)$ denotes the similarity between images u and v. Intuitively, $f(S)$ captures a trade-off between coverage and diversity: the first term rewards similarity between the selected images and all images, while the second term penalizes redundancy within S. As shown by Mirzasoleiman et al. [36], $f(\cdot)$ is a non-monotone submodular function.

Following Cui et al. [13] and [18, 23, 25], we use a subset of the CIFAR-10 dataset [29] containing images from the *Airplane*, *Automobile*, and *Bird* classes, which consists of tens of thousands of 32×32 color images, and define $\mathrm{Sim}(u, v)$ as the inner product of the 3,072-dimensional pixel

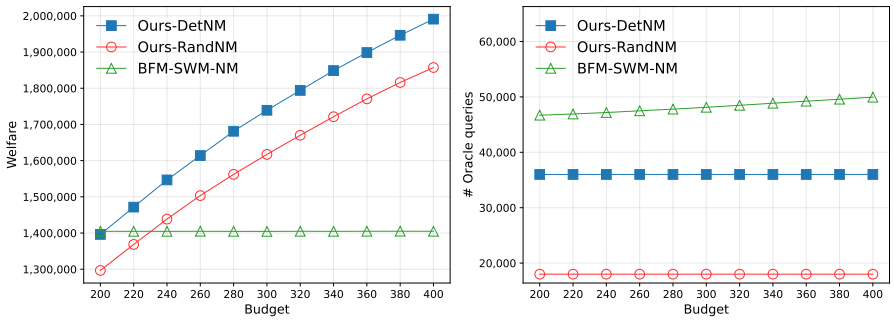

Figure 2: Comparison of OURS-DETNM, OURS-RANDNM, and BFM-SWM-NM on the crowdsourcing application in terms of social welfare and the number of value-oracle queries.

vectors of images u and v. Again following Cui et al. [13] and [23, 36], the cost of each image is set proportional to the standard deviation of its pixel intensities and normalized so that the average cost is 0.1. Consequently, images exhibiting greater variation in pixel intensities are assigned higher costs, whereas images with more uniform pixel intensities are assigned lower costs.

With the objective of maximizing $f(S) - c(S)$, the above crowdsourcing application falls within our model of budget-feasible submodular welfare maximization under non-monotone submodular valuations, a setting considered only in our work and in Cui et al. [13]. We therefore compare the following mechanisms: (1) our Algorithm 1 with input parameters chosen as specified in Theorem 5.2, denoted by **OURS-RANDNM**; (2) our Algorithm 1 with input parameters chosen as specified in Theorem 5.3, denoted by **OURS-DETNM**; and (3) the auction mechanism of Cui et al. [13], denoted by **BFM-SWM-NM**, with parameters optimized for non-monotone submodular valuations (cf. Theorem 4.8 therein) and $\epsilon = 0.1$, following their experimental setting.

Experimental Results. Figure 2 compares the implemented mechanisms in terms of social welfare and the number of value-oracle queries. Since OURS-RANDNM is randomized, its results are averaged over 50 runs; the standard deviation ranges from 0.0066×10^6 to 0.0084×10^6, which is below 0.7% of the smallest welfare observed. The results show that both OURS-RANDNM and OURS-DETNM consistently outperform BFM-SWM-NM in most cases, achieving substantially higher social welfare with significantly fewer value-oracle queries. In particular, relative to BFM-SWM-NM, OURS-DETNM and OURS-RANDNM improve social welfare by up to 41.7% and 32.2%, respectively, while reducing the number of value-oracle queries by up to 27.9% and 64.0%, respectively. On average, OURS-DETNM and OURS-RANDNM achieve 22.5% and 14.0% higher social welfare than BFM-SWM-NM, respectively, while using 25.3% and 62.7% fewer value-oracle queries, as shown by Table 2.

Table 2: Average performance of the implemented mechanisms across all tested points in Figure 2.

Mechanism	Welfare ↑	Oracle queries ↓
OURS-DETNM	**1,720,667**	36,002
OURS-RANDNM	1,601,964	**18,002**
BFM-SWM-NM	1,404,784	48,209

We also observe that OURS-RANDNM generally yields lower social welfare than OURS-DETNM, although it attains the best approximation ratio and requires the fewest oracle queries. One possible explanation is that the approximation guarantee of OURS-RANDNM holds only in expectation. In addition, since OURS-RANDNM is randomized, each seller accepting the offered price is discarded with probability approximately 0.63, which may remove sellers with high-quality images. Despite this, OURS-RANDNM still outperforms BFM-SWM-NM.

In summary, the above experimental results indicate that our Algorithm 1 significantly outperforms the mechanism of Cui et al. [13] under non-monotone submodular valuations in terms of both social welfare and query complexity.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction accurately reflect the proposed mechanism, its theoretical guarantees, and the experimental evaluation.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Yes. In the conclusion, the paper notes that although it establishes improved approximation guarantees for both monotone and non-monotone submodular valuations, it remains unclear whether these these guarantees are tight.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All theorems and lemmas are stated with full assumptions and complete proofs. Due to space limits, most proofs are deferred to the appendix.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Section 5 and Appendix C describe the proposed mechanism, baselines, datasets, and parameter settings in sufficient detail to reproduce the experimental results.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The paper does not provide open access to code at this time. However, all algorithms are described in sufficient detail to enable faithful reimplementation, and all datasets used (SNAP social networks and CIFAR-10) are publicly available.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

Answer: [Yes]

Justification: Section 5 and Appendix C specify the datasets, valuation functions, budget ranges, parameter settings of all mechanisms (taken from Theorems 4.1–4.3 and from the published settings of the baselines), and the number of random seeds.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: For the randomized mechanism OURS-RANDNM, we report the mean over 50 independent seeds per budget, and the sample standard deviation across seeds is reported in Appendix C. All other mechanisms are deterministic and produce a single outcome per instance.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Section 5 reports the compute setup used in our experiments (CPU-only Apple M4 Pro, 24 GB memory). Our experiments compare the time complexity of the evaluated methods to characterize their computational cost.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research conforms to the NeurIPS Code of Ethics. It is theoretical and algorithmic, uses only publicly available datasets, and does not involve human subjects.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [No]

Justification: The paper mentions applications such as crowdsourcing, influence maximization, and federated learning, but does not include a dedicated discussion of broader societal impacts. Such impacts are indirect and depend on downstream deployment contexts.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: The paper does not release pre-trained models, generative models, or scraped datasets, and poses no such risks.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All datasets used (SNAP social networks and CIFAR-10) are publicly available for research, and all prior algorithms used as baselines are properly cited.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.

- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not release new datasets or models.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or human subjects; the "crowdsourcing" application is an algorithmic benchmark on CIFAR-10 with synthetic costs.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve research with human subjects.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLMs were not used as an important, original, or non-standard component of the core methodology.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.