

Mechanism Design for Daily Deals

Binyi Chen¹, Tao Qin², Tie-Yan Liu²

¹University of California, Santa Barbara ²Microsoft Research
binyichen@cs.ucsb.edu {taoqin, tyliu}@microsoft.com

ABSTRACT

Daily deals are very popular in today’s e-commerce. In this work, we study the problem of mechanism design for a daily deal website to maximize its revenue and obtain the following results. (1) For the Bayesian setting, we first design a revenue-optimal incentive-compatible (IC) mechanism with pseudo-polynomial time complexity. Considering the high computational complexity of the mechanism, we then develop a greedy mechanism that is much more computationally efficient yet maintains a constant competitive ratio regarding the Bayesian optimal revenue in expectation. (2) For the prior-independent setting, we first propose a randomized IC mechanism with a pseudo-polynomial time complexity that can achieve a constant competitive ratio. Then, by leveraging the greedy mechanism designed for the Bayesian setting, we come up with a new mechanism that can achieve a good tradeoff between computational efficiency and competitive ratio. After that, we discuss the robustness issue regarding the two mechanisms (i.e., they both use the trick of random partition and may perform badly for the worst-case partition) and propose an effective way to guarantee a constant competitive ratio even for the worst-case partition.

Categories and Subject Descriptors

J.4 [Computer Applications]: Social and Behavioral Sciences–Economics

General Terms

Economics, Theory

Keywords

Mechanism Design, Daily Deals

1. INTRODUCTION

Daily deal services (e.g., Groupon) have become very popular nowadays. In daily deal services, a merchant demands and pays for a chunk of user impressions.¹ The demand of

¹In daily deal services, merchants usually target at conver-

a merchant can be specified by a pair (l, u) , where l and u denote the minimum and maximum numbers of user impressions that she would like to buy. The daily deal website can collect revenue from the merchant if the number of user impressions allocated to him/her (denoted as x) satisfies $l \leq x$: if $x < l$, the merchant will pay nothing because the deal is not tipped on; if $x > u$, the merchant will not pay extra money for the user impressions beyond u .

Daily deal services have been studied from many different perspectives, such as merchants’ reputation [7] and profitability [11, 12], and customers’ behaviors [6, 27, 21, 26]. Different from those work, we focus on the problem of revenue maximization by means of designing auction mechanisms.

1.1 Daily Deal Auctions

We define the Daily Deal Auctions (DDAs) as follows. Let N be the supply from the daily deal website, i.e., there are N web users visiting the website during a specific time period. Denote K as the maximum number of slots where the website can show deals to each web user. We follow the convention in previous work [13, 1, 19] and assume the qualities of the slots to be different from each other (e.g., a deal shown at the top slot is more likely to attract users’ attention). Specifically, we use λ_k to denote the quality of slot k and assume that the slots are numbered in the descending order of their qualities: $1 \geq \lambda_1 > \lambda_2 > \dots > \lambda_K \geq 0$.

If a deal is shown to x users at slot k , we say that the deal has obtained $x\lambda_k$ effective impressions. We define $N_k = N\lambda_k$ as the number of effective impressions of slot k in the entire time period (another way to understand N_k is to regard it as the expected number of users who have paid attention to the k -th slot). For simplicity and without much loss of accuracy, we assume N_k to be an integer. It is clear that $N_1 > N_2 > \dots > N_K$.

Let M be the number of candidate deals participating in the auction. Each deal i is characterized by a tuple (l_i, u_i, v_i) , where l_i and u_i denote the minimum and maximum number of effective impressions demanded by the deal, and v_i denotes the per-effective-impression value. We focus on the single-parameter setting, in which a merchant may strategically misreport her value v_i through a bid b_i .

sions/purchases, which can be easily transformed to impressions. For example, 10000 impressions is equivalent to 100 conversions if the conversion rate of the deal is 0.01. We choose to use the term *impression* throughout the paper because it is what the daily deal website can allocate in practice.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.

Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Upon receiving the bids $\{b_i\}_{i \in [M]}$ from all the merchants, the website determines how to allocate the slots and how to charge merchants by using a mechanism. Let x_i be the number of effective impressions allocated to bidder i and p_i be the payment of bidder i . Then the utility of bidder i is

$$U_i = 1\{l_i \leq x_i\}(\min\{x_i, u_i\}v_i - p_i),$$

where $1\{\cdot\}$ is the indicator function, and the revenue of the auctioneer is $\sum_{i \in [m]} p_i$. Note that the utility of any bidder should be non-negative for *individual rationality*.

A feasible allocation for DDAs must satisfy two conditions: (1) no more than one deal can be assigned to a slot for any user; (2) no deal can be assigned to more than one slot for any user. These feasibility conditions can be mathematically expressed by the following majorization constraints [19]:

$$\sum_{j=1}^i x_{[j]} \leq \sum_{j=1}^i N_j, \forall i \in [M],$$

where $x_{[j]}$ denotes the j -th largest element in $\{x_i\}_{i \in [M]}$, and $N_i = 0, \forall K < i \leq M$.

The goal of our work is to design computationally efficient mechanisms for DDAs, which are incentive compatible to the bidders and can (approximately) maximize the revenue of the auctioneer.

1.2 Our Results

Our main results can be summarized as follows.

We start with the *Bayesian setting*. (1) We first propose an incentive compatible (IC) mechanism \mathcal{M}_1 with a pseudo-polynomial time complexity, which can achieve the optimal expected revenue. This mechanism leverages the dynamic program proposed in [19] for the allocation, and uses a novel algorithm to calculate the payment. (2) Considering the high complexity of \mathcal{M}_1 , we further propose a greedy mechanism \mathcal{M}_2 that is highly computationally efficient (e.g., $O(M^3)$) and can achieve a competitive ratio of $1/4$.

We then consider the *prior-independent setting*. (1) Based on \mathcal{M}_1 , we design a randomized IC mechanism \mathcal{M}_3 with a pseudo-polynomial time complexity and a competitive ratio of $\frac{\beta-1}{\beta}$, where $\beta > 1$ is a fixed constant. (2) To reduce the computational complexity of \mathcal{M}_3 , we develop a greedy mechanism \mathcal{M}_4 with complexity is $O(M^3)$ and a competitive ratio of $\frac{\beta-1}{4\beta}$. (3) Considering that the above two mechanisms rely on random partition of the bidders and only have performance guarantee in the expected case, we then discuss how to improve their robustness, i.e., how to guarantee their performances even for the worst-case partition. As a showcase, we develop a mechanism \mathcal{M}_5 that can achieve a competitive ratio of $\frac{3}{7}$ and a robust ratio (see Definition 10) of $\frac{1}{4}$.

2. RELATED WORK

Daily deal services have been studied from many different perspectives. [7] studies how daily deal sites affecting the reputation of a business using evidences from Yelp reviews. [6] investigates hypotheses such as whether daily deal subscribers are more critical than their peers. [11] tries to answer the the question whether group-buying deals would be profitable for businesses. [12] finds that offering vouchers is more profitable for merchants which are patient or relatively unknown, and for merchants with low marginal costs.

There are several papers [27, 21, 26] studying consumer purchase/repurchase behaviors towards daily deals. It is easy to see that the focuses of existing pieces of work are very different from ours. We stand on the position of a daily deal site and focus on the problem of revenue maximization by means of designing auction mechanisms.

A closely related work is [8], in which the authors also design auction mechanisms for daily deals. The difference is that they target at the maximization of a combination of the revenue of the auctioneer (website), welfare of the bidders (merchants), and the utility of the consumers, while we focus on the revenue of the website. Besides, they ignore the dynamic demand of merchants (i.e., from l_i to u_i) and do not consider the impression allocation problem.

The impression allocation problem in our work is related to the classical knapsack problem [9], which can be viewed as our special case by setting $l_i = u_i$ and $K = 1$. While some recent work on stochastic knapsack problem [10] consider items with stochastic volumes, the allocation of an item is binary (i.e., pack it or not). Our work is different in the sense that the allocation of a deal is not binary (zero or lower bounded by l_i and upper bounded by u_i). Our work is also highly related to [19] and [20], except that they ignore the incentive issues and assume that bidders truthfully report their valuations.

From the perspective of auction design, DDAs are related to knapsack auctions as studied in [2]. In fact, our problem will degenerate to the problem under their investigations by (1) setting $l_i = u_i, K = 1$ and (2) assuming bidders' private value are deterministic and fixed. Recently, [25] studies knapsack auctions in the Bayesian setting and obtain several approximation results, yet it is still a special case of our setting ($l_i = u_i$ and $K = 1$). DDAs are also related to multi-unit auctions [22, 16]. The difference is that the allocation in DDAs have a lower (l_i) and an upper (u_i) constraint.

[5] and [4] are the early pieces of work that consider payment computation in truthful mechanisms. Both of them and our work face the hinder that the computation involves the integral of a step function with possibly exponential breakpoints. Our method is different from them and we solve this issue in the specific setting of daily deal auctions.

The concept of competitive ratio comes from [18], which evaluates the performance of a mechanism by comparing its expected revenue against the Bayesian optimal expected revenue. The concept of "prior-independent setting" also comes from [18]. Our designed mechanisms for the prior-independent setting are inspired by [14], [15] and [3] in which bidders are partitioned into several sub markets to estimate the value distribution. Again, different from these pieces of work on digital goods, we consider bidders with dynamic demands (ranging from l_i to u_i).

3. PRELIMINARIES

3.1 Settings and Notations

We mainly investigate on two settings for DDAs, and use *competitive ratio* to measure the expected revenue that an auction mechanism can achieve.

In the *Bayesian setting*, it is assumed that the valuation v_i of bidder i is independently sampled from a distribution F_i , which is known to both the auctioneer and all the bidders. In this setting, the competitive ratio of a mechanism \mathcal{M} is defined as $\frac{Rev(\mathcal{M})}{OPT}$, where OPT denotes the optimal expected

revenue of an incentive compatible mechanism and $Rev(\mathcal{M})$ is the expected revenue of \mathcal{M} . Note that both expectations are taken over the joint distribution $F_1 \times F_2 \times \dots \times F_M$.

In the *prior-independent setting*, it is assumed that the valuations of all the bidders are independently sampled from the same distribution F , but this distribution is unknown to the auctioneer. In this setting, the competitive ratio of a mechanism \mathcal{M} is defined as $\inf_F \frac{Rev(\mathcal{M}, F)}{OPT(F)}$, where $OPT(F)$ denotes the optimal expected revenue of an incentive compatible mechanism and $Rev(\mathcal{M}, F)$ is the expected revenue of \mathcal{M} . Both expectations are taken with respect to the distribution F .

In addition, throughout the paper, we use v_{-i} to denote the valuation profile of all the other bidders except i , use $x_i(v_i) = \mathbf{E}_{v_{-i}}[x_i(v_i, v_{-i})]$ to denote the expected number of effective impressions allocated to bidder i when her valuation is v_i , use $p_i(v_i)$ to denote the expected payment, and use $u_i(v_i) = v_i x_i - p_i$ to denote the expected utility. Since all the mechanisms designed in this paper are Bayesian incentive compatible, we will interchangeably use the valuation profile v and bid profile b , and so for v_i and b_i .

3.2 Myerson's Lemma

As mentioned in the introduction, we are concerned with the *single-parameter* auctions, in which each bidder's private information can be expressed by a one-dimensional value v_i . We can exploit the classical *Myerson Lemma* to design incentive compatible mechanisms for single-parameter auctions:

LEMMA 1. *The sufficient conditions for a mechanism to be Bayesian incentive compatible are as follows.*

1. *Allocation monotonicity: for all i and $v_i > v'_i$, $x_i(v_i) \geq x_i(v'_i)$.*
2. *Payment identity: for all i and v_i , $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(v) dv$.*

Furthermore, we can leverage the following lemma to design incentive-compatible revenue-optimal mechanisms.

LEMMA 2. [24] *In a Bayesian incentive compatible mechanism with allocation rule $\mathbf{x}(\cdot)$, the expected payment of agent i satisfies $\mathbf{E}_{\mathbf{v}}[p_i(v)] = \mathbf{E}_{\mathbf{v}}[\phi_i(v_i)x_i(v)]$, where we denote the virtual value $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$, and f_i and F_i are the probability density function and cumulative distribution function separately for bidder i 's valuation.*

Following the common practice [24, 23] in mechanism design, we assume that the valuation distributions (all F_i 's in the Bayesian setting and F in the prior-independent setting) are regular (i.e., the virtual value $\phi_i(v_i)$ is monotone with respect to v_i).

4. BAYESIAN SETTING

In this section, we design mechanisms for DDAs in the Bayesian setting. The goal is to (approximately) maximize the revenue through incentive compatible mechanisms.

4.1 Bayesian Optimal Mechanism

We design a mechanism \mathcal{M}_1 that leverages the dynamic program proposed for the Chunked Allocation Problem in [19], together with a tie-breaking trick, to find the virtual

surplus maximizing allocation, and charges bidders according to Eqn. (2). The details of the allocation and payment rules are discussed in Section 4.1.1 and 4.1.2 respectively. For ease of reference, we denote the dynamic program proposed in [19] as DP-CAP.

As will be shown in Section 4.1.3, \mathcal{M}_1 is Bayesian incentive compatible. In addition, because the allocation rule maximizes the virtual surplus, this mechanism is revenue optimal according to Lemma 2.

Algorithm 1 A Bayesian Optimal Mechanism (\mathcal{M}_1)

Input: $\{N_k\}_{k \in [K]}, \{b_i, l_i, u_i, F_i\}_{i \in [M]}$

Output: x_i , the number of effective impressions allocated to bidder i , and p_i , the payment of bidder i .

- 1: Compute bidders' virtual values, and remove bidders with negative virtual values.
 - 2: Find an allocation that maximizes virtual surplus using DP-CAP plus the tie-breaking trick described in Section 4.1.1. Let x_i be the number of effective impressions allocated to bidder i .
 - 3: **for** each bidder $i \in [M]$ **do**
 - 4: **if** $x_i = 0$ **then**
 - 5: $p_i \leftarrow 0$.
 - 6: **else**
 - 7: For all $a \in \{0, l_i, l_i + 1, \dots, x_i\}$, compute $\Phi_{-i}(a)$, the maximum virtual surplus generated from all other bidders except i under the condition that i is allocated with exactly a effective impressions.
 - 8: Construct a convex hull by intersecting $x_i - l_i + 2$ half planes: $f_a(y) = ay + \Phi_{-i}(a) \geq 0, \forall a \in \{0, l_i, l_i + 1, \dots, x_i\}$. Each corner point of the convex hull can be represented by a pair, the first element of which indicates the virtual value of bidder i , and the second element of which is the number of effective impressions allocated to bidder i . We can get a set of pairs: $(0, 0), (t_1, k_1), \dots, (t_q, k_q), (\phi_i(b_i), x_i)$, where $0 < t_1 < t_2 < \dots < t_q < \phi_i(b_i)$.
 - 9: $p_i \leftarrow x_i \cdot z_q - \sum_{j=1}^q k_j \cdot (z_j - z_{j-1})$, where $z_j = \phi_i^{-1}(t_j)$.
 - 10: **end if**
 - 11: **end for**
 - 12: **return** $\{x_i, p_i\}_{i \in [M]}$
-

4.1.1 The Allocation Rule

While the allocation rule of \mathcal{M}_1 looks quite simple since it leverages an existing algorithm to find an optimal allocation, we would like to point out an unobvious pitfall here. Please note that the direct application of DP-CAP cannot guarantee the monotonicity of the allocation rule, when the optimal allocations are not unique. For example, suppose $\phi_i(b_i) = \phi_i(b'_i)$ for $b'_i < b_i$. Since there may be multiple allocations that achieve the optimal virtual surplus, it is possible that for bidder i , the allocation that she obtains when bidding b_i is less than the allocation when bidding b'_i .

To address this issue, we need to break the ties so as to ensure the unique output of DP-CAP, and therefore ensure the monotonicity of the allocation rule of \mathcal{M}_1 . Here we take the simple case of $K = 1$ (i.e., there is only one slot) as an example for illustration. For the case of multiple slots ($K > 1$), the same tie-breaking trick also applies.

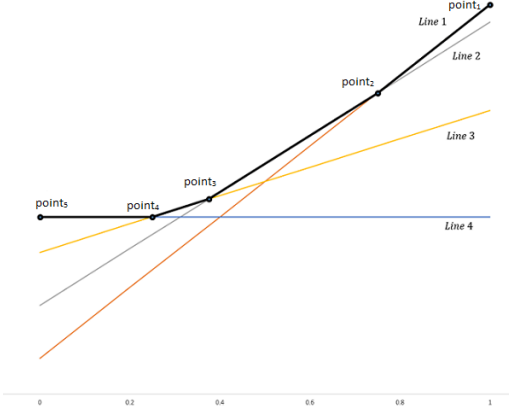


Figure 1: Computation of the Payment Rule

In this simple case, the transition equation of DP-CAP is

$$VS(i, n) = \max_{x_i} (VS(i-1, n-x_i) + x_i \phi_i) \quad x_i \in [l_i, u_i] \cup \{0\}, \quad (1)$$

where $VS(i, n)$ represents the maximum virtual surplus from n effective impressions and bidders $\{1, 2, \dots, i\}$.

The key of ensuring the unique output of DP-CAP is as follows. When multiple x_i 's make $VS(i-1, n-x_i) + x_i \phi_i$ maximal, we simply choose the minimum one among them. The optimal allocation is the one associated with $\max(VS(M, n))$ for $n \in [N]$. If there are multiple n 's that achieve the maximum, we choose the minimum n .

4.1.2 The Payment Rule

According to Lemma 1, the payment rule for a Bayesian incentive compatible auction is:

$$p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(v) dv. \quad (2)$$

At first glance, given the above result, it seems trivial to design the payment rule for our problem. For example, when x_i takes binary values (e.g., win or lose), similar to the classical VCG payment [17], we can simply exploit the approach proposed in [18] to calculate the payment: let $p'_i = OPT(v_{-i}) - OPT_{-i}(v)$, where $OPT(v_{-i})$ is the optimal virtual surplus when i has been removed and $OPT_{-i}(v)$ is the original optimal virtual surplus minus i 's virtual surplus, and charge i with the price $p_i = \phi_i^{-1}(p'_i)$.

However, we note that in daily deal auctions, x_i does not necessarily take binary values. In fact, it can take any value from $\{0, l_i, l_i + 1, \dots, \min(u_i, \text{allocated } N_j)\}$. In this case, the design of the payment rule becomes harder, and we need to explore the structural properties of the allocation rule. Since the allocation rule is monotone, fixing the bids of all the other bidders, x_i is a piecewise constant function of b_i . For simplicity, we call those bids corresponding to the change of x_i as breakpoints. To compute $p_i(v_i)$ in Eqn. (2), we need to know all the breakpoints of bidder i . As shown in Steps 7-9 of \mathcal{M}_1 , we can find all the breakpoints by calling DP-CAP for at most $x_i(v_i) - l_i + 2$ times, and thus the total complexity of payment computation is $\text{Poly}(N, M, K)$. The remaining part of this subsection explains how Steps 7-9 work.

Consider the payment of bidder i when she is allocated with $x_i > 0$ effective impressions. Let $\Phi_{-i}(a)$ denotes the part of the maximum virtual surplus generated from all the

other bidders under the condition that a fixed number of effective impressions (denoted as a) are allocated to bidder i . $\Phi_{-i}(a)$ can be easily calculated by calling DP-CAP and setting all its states with $x_i \neq a$ to negative infinity. Then for each allocation $a \in \{0, l_i, l_i + 1, \dots, x_i(v_i)\}$, the maximum virtual surplus (conditioned on that i gets a effective impressions) can be expressed as $a \cdot \phi_i + \Phi_{-i}(a)$, which is a linear function of ϕ_i with slope a and bias $\Phi_{-i}(a)$. After calling DP-CAP for $x_i(v_i) - l_i + 2$ times, we get $x_i(v_i) - l_i + 2$ lines with slopes $\{0, l_i, l_i + 1, \dots, x_i(v_i)\}$ and biases $\{\Phi_{-i}(0) \geq \Phi_{-i}(l_i) \geq \Phi_{-i}(l_i + 1) \geq \dots, \geq \Phi_{-i}(x_i(v_i))\}$.

As shown in Figure 4.1.2, the horizontal axis is the virtual value of bidder i and the vertical axis is the corresponding virtual surplus. First we set Cur to be the line with the largest slope (in Figure 4.1.2 it is line 1) and set (cx, cy) to be the point $(\phi_i(v_i), \Phi_{-i}(x_i) + x_i \cdot \phi_i(v_i))$. In each step, we can calculate the intersection points between the line Cur and other lines in $O(N)$ time, and choose the upmost point on the left side of the point (cx, cy) . We use a pair to denote this upmost point: the first element of the pair is the virtual value of bidder i at the point and the second one is the number of effective impressions allocated to i . We then update (cx, cy) to be this point and update Cur to be the corresponding intersecting line. We terminate the process when $cx \leq 0$. Since each line can become Cur in at most one step, the procedure will terminate in $O(N)$ steps. Finally, we find those points (and their associated pairs) that build the convex hull of the lines (in Figure 4.1.2, they are $point_1, \dots, point_5$). Assume those pairs are $(0, 0), (t_1, k_1), \dots, (t_q, k_q), (\phi_i(v_i), x_i(v_i))$, where $0 < t_1 < t_2 < \dots < t_q < \phi_i(v_i)$. We set $t_0 = 0$. For all $j \in \{0, \dots, q\}$, we set $z_j = \phi_i^{-1}(t_j)$, and compute the payment of bidder i as below:

$$p_i(v_i) = x_i(v_i) \cdot z_q - \sum_{j=1}^q k_j \cdot (z_j - z_{j-1}).$$

4.1.3 Incentive Compatibility

THEOREM 3. \mathcal{M}_1 is Bayesian incentive compatible.

PROOF. Lemma 1 has given the sufficient conditions for a mechanism to be Bayesian incentive compatible. It is easy to see that the payment rule of \mathcal{M}_1 satisfies the second condition. Therefore we only need to prove that \mathcal{M}_1 satisfies the first condition: its allocation rule is monotone.

Let $x^*(b)$ denote the unique optimal allocation vector outputted by the dynamic program for the bid profile b , and $x_i^*(b)$ denote the number of effective impressions allocated to bidder i . Without loss of generality, we assume $x_i^*(b) > 0$. Let $\Phi(x, b)$ denote the virtual surplus when the bid profile is b and allocation is x .

Now we prove the monotonicity of $x^*(b)$ by contradiction. If $x^*(b)$ is not monotone, there exist some $i, b_{-i}, b_i < b'_i$ such that $x_i^*(b) > x_i^*(b')$, where $b = (b_i, b_{-i})$ and $b' = (b'_i, b_{-i})$.

If $\phi_i(b_i) = \phi_i(b'_i)$ the allocation of i will not change, thus we can assume $\phi_i(b_i) < \phi_i(b'_i)$. Since $x_i^*(b) > 0$ is the optimal allocation for b and $x^*(b')$ is the optimal allocation for b' , we have $\Phi(x^*(b), b) \geq \Phi(x^*(b'), b)$ and $\Phi(x^*(b'), b') \geq \Phi(x^*(b), b')$. Then we get

$$\begin{aligned} \Phi(x^*(b), b) - \Phi(x^*(b), b') &\geq \Phi(x^*(b'), b) - \Phi(x^*(b'), b') \\ \Rightarrow x_i^*(b)(\phi_i(b_i) - \phi_i(b'_i)) &\geq x_i^*(b')(\phi_i(b_i) - \phi_i(b'_i)) \\ \Rightarrow x_i^*(b) &\leq x_i^*(b') \end{aligned}$$

Thus we arrive at a contradiction. \square

4.2 A Greedy Mechanism

Although \mathcal{M}_1 is revenue optimal, it suffers from the huge computational complexity. Because it relies on DP-CAP to find the virtual surplus maximizing allocation, its running time is quadratic to N . This is not practically feasible because a typical publisher for daily deal auctions may have millions (even hundreds of millions) of user impressions every day. In this subsection, we propose a simple greedy mechanism \mathcal{M}_2 (see Algorithm 2), which is much more computationally efficient and has a revenue guarantee.

Algorithm 2 A $\frac{1}{4}$ -Approximation Mechanism (\mathcal{M}_2)

- 1: Calculate bidders' virtual value. Remove those bidders with negative virtual values. Let $S = \{1, \dots, t\}$ denote the remaining t bidders.
 - 2: Sort bidders according to their virtual bidding values, w.l.o.g $\phi_1(b_1) \geq \phi_2(b_2) \geq \dots \geq \phi_t(b_t)$.
 - 3: Partition S into \tilde{K} parts: $\{1, \dots, t_1\}, \{t_1 + 1, \dots, t_2\}, \dots, \{t_{\tilde{K}-1} + 1, \dots, t_{\tilde{K}}\}, \{t_{\tilde{K}} + 1, \dots, t\}$ ($\tilde{K} < K$ if and only if $t_{\tilde{K}} = t$) as follows:
$$N_1 + u_{t_1} \geq u_1 + u_2 + \dots + u_{t_1} > N_1$$

$$N_2 + u_{t_2} \geq u_{t_1+1} + u_{t_1+2} + \dots + u_{t_2} > N_2$$

$$\dots \geq \dots + \dots + \dots > \dots$$

$$N_{\tilde{K}} + u_{t_{\tilde{K}}} \geq u_{t_{\tilde{K}-1}+1} + u_{t_{\tilde{K}-1}+2} + \dots + u_{t_{\tilde{K}}} > N_{\tilde{K}}$$
 - 4: **for** each slot $i \leq \tilde{K}$ **do**
 - 5: With probability $1/2$, each bidder $j \in \{t_{i-1}+1, \dots, t_i-1\}$ gets u_j effective impressions on slot i with payment computed according to Case 1 in Section 4.2.1 (here we set $t_0 = 0$).
 - 6: With probability $1/2$, bidder t_i gets $\min(N_i, u_{t_i})$ effective impressions on slot i , with payment computed according to Cases 2 and 3 in Section 4.2.1.
 - 7: **end for**
 - 8: **return** $\{x_i, p_i\}_{i \in [M]}$.
-

4.2.1 Payment Computation

While the allocation rule is simple and clearly described in Algorithm 2, the payment rule is much more complicated. To compute the payments of all the winning bidders, we need to consider the following three cases.

Case 1: For bidder $j \in \{t_{i-1} + 1, \dots, t_i - 1\}$, if j 's expected allocation is non-zero, we calculate a set of critical bids $\{b_j(0), b_j(s), b_j(s+1), \dots, b_j(\tilde{K})\}$ for him/her, where $s = \arg \min_k (N_k < u_j)$, $b_j(0)$ is the minimum bid for bidder j to win $\frac{u_j}{2}$ effective impressions in expectation when others' bids are fixed, and $b_j(k)$ is the minimum bid for him/her to win $\frac{N_k}{2}$ effective impressions in expectation when others' bids are fixed. It is easy to verify that $b_j(0) > b_j(s) > b_j(s+1) > \dots > b_j(\tilde{K})$. $b_j(k)$ can be calculated by enumerating the bid from $\{\phi_j^{-1}(\phi_m(b_m))\}_{m=1}^t$ and finding the new slot that bidder j is allocated to. The complexity of calculating the set $b_j(\cdot)$ for j is $O(M^2)$. Then the payment for bidder j is

$$p_j = b_j(0) \cdot u_j - (b_j(0) - b_j(s)) \cdot N_s - \sum_{l=s+1}^{\tilde{K}} (b_j(l-1) - b_j(l)) \cdot N_l.$$

Case 2: For bidder t_i , if t_i 's expected allocation is non-zero and $\min(N_i, u_{t_i}) = u_{t_i}$, the payment can be computed in the same way as in the first case.

Case 3: For bidder t_i , if t_i 's expected allocation is non-zero and $\min(N_i, u_{t_i}) = N_i$, we can compute the set of critical bids in the same way as in the first case, and then compute the payment as below.

$$p_{t_i} = b_{t_i}(i) \cdot N_i - \sum_{l=i+1}^{\tilde{K}} (b_{t_i}(l-1) - b_{t_i}(l)) \cdot N_l.$$

Since at most M bidders can win effective impressions and the complexity of payment computation for each bidder is $O(M^2)$, the complexity of \mathcal{M}_2 is $O(M^3)$, independent of the number of user impressions. In this regard, \mathcal{M}_2 is much more computationally efficient than \mathcal{M}_1 , because the number of merchant (e.g., a typical daily deal website usually have hundreds of candidate deals at the same time.) is much smaller than that of the user impressions.

4.2.2 Theoretical Analysis

It is not difficult to verify that the allocation rule of \mathcal{M}_2 is monotone and the payment rule satisfies the condition in Lemma 1. As a consequence, \mathcal{M}_2 is Bayesian incentive compatible. Here we mainly investigate the competitive ratio of \mathcal{M}_2 . We first prove two lemmas and then give the main theorem regarding the competitive ratio.

LEMMA 4. Denote $OPT(\mathbf{N}, b)$ as the optimal virtual surplus for a bid profile b under the supply vector \mathbf{N} . For \mathcal{M}_2 we have

$$\sum_{k=1}^{\tilde{K}} \left[\sum_{j=t_{k-1}+1}^{t_k-1} \phi_j(b_j) \cdot u_j + \phi_{t_k}(b_{t_k}) \cdot N_k \right] \geq OPT(\mathbf{N}, b).$$

PROOF. Recall that the bidders are numbered according to the non-increasing order of their virtual values.

Let \hat{x} be an allocation vector defined as below,

$$\hat{x} = (u_1, u_2, \dots, u_{t_1-1}, N_1, u_{t_1+1}, \dots, u_{t_2-1}, N_2, u_{t_2+1}, \dots, u_{t_{\tilde{K}-1}}, N_{\tilde{K}}, 0, 0, \dots, 0).$$

Note that \hat{x} over allocates the effective impressions and therefore is not feasible. It will be used to upper bound the virtual surplus of the optimal allocation vector.

We use \hat{x}_i to denote the i -th dimension of \hat{x} , i.e., the number of effective impressions allocated to bidder i in allocation vector \hat{x} .

Let x^* be the allocation profile given by the Bayesian optimal auction (i.e., mechanism \mathcal{M}_1). Let ϕ be the virtual value vector: $\phi = (\phi_1(b_1), \dots, \phi_t(b_t))$. The left side of the inequality in the lemma therefore can be written as $\langle \phi, \hat{x} \rangle$, and the right side can be written as $\langle \phi, x^* \rangle$.² In addition, we use $\vec{1}$ to denote the t -dimensional identity vector.

First, we prove $\langle \vec{1}, \hat{x} \rangle \geq \langle \vec{1}, x^* \rangle$ by considering two possible cases with respect to \tilde{K} .

Case 1: $\tilde{K} < K$. In this case, we have $t_{\tilde{K}} = t$ and $\hat{x} = (u_1, u_2, \dots, u_{t_1-1}, N_1, u_{t_1+1}, \dots, u_{t_2-1}, N_2, u_{t_2+1}, \dots, u_{t_{\tilde{K}-1}}, N_{\tilde{K}})$. It is clear $\forall i, u_i \geq x_i^*$. According to the majorization constraints, we have

$$\sum_{j=1}^{\tilde{K}} x_{t_j}^* \leq \sum_{j=1}^{\tilde{K}} x_{[j]}^* \leq \sum_{j=1}^{\tilde{K}} N_j.$$

Therefore we obtain $\langle \vec{1}, \hat{x} \rangle \geq \langle \vec{1}, x^* \rangle$.

Case 2: $\tilde{K} = K$. In this case, we have $\langle \vec{1}, \hat{x} \rangle \geq \sum_{j=1}^K N_K \geq \langle \vec{1}, x^* \rangle$.

²Here we use $\langle \cdot, \cdot \rangle$ to denote the inner product of two vectors.

Second, we prove $\langle \phi, \hat{x} \rangle \geq \langle \phi, x^* \rangle$. For ease of reference, we define $r_0 = \langle \phi, \hat{x} \rangle$. The basic idea of the proof is to construct another allocation vector x , which leads to a virtual surplus smaller than r_0 but larger than $\langle \phi, x^* \rangle$. We initialize $x = \hat{x}$, and then modify its elements one by one in the following steps.

(1) Considering $\min(N_1, u_1) \geq x_1^*$, in the first step, we decrease x_1 from \hat{x}_1 to x_1^* and increase x_2 from \hat{x}_2 to $\hat{x}_2 + \hat{x}_1 - x_1^*$. Define $r_1 = \phi \cdot x$. Then we have $r_1 \leq r_0$ because $\phi_1(b_1) \geq \phi_2(b_2)$. Now we show that after the update, $x_2 \geq x_2^*$:

- If $\hat{x}_2 = u_2$, we have $x_2 \geq \hat{x}_2 \geq x_2^*$.
- Otherwise, $\hat{x}_2 = N_p$. Considering the majorization constraints and $u_i \geq x_i^*$, we get

$$\begin{aligned} x_2 &= N_p + \hat{x}_1 - x_1^* = \sum_{j=1}^p N_j + [\hat{x}_1 \neq N_1]u_1 - x_1^* \\ &\geq x_1^* + x_2^* - x_1^* = x_2^*. \end{aligned}$$

(2) In the i -th step ($i < t_{\hat{K}}$), we have $x_i \geq x_i^*$. We increase x_{i+1} to $\hat{x}_{i+1} + x_i - x_i^*$ and decrease x_i to x_i^* . We define $r_i = \langle \phi, x \rangle$ and have $r_i \leq r_{i-1}$. Note that the allocation vector x has been updated and therefore r_i 's are different from each other. Similarly, we can verify that $x_{i+1} \geq x_{i+1}^*$.

By going through the above steps, we modify all the elements in x and obtain a sequence of r_i . Finally, we have the lemma proved as follows.

$$\begin{aligned} \langle \phi, \hat{x} \rangle &= r_0 \geq r_{t_{\hat{K}}-1} \\ &= \sum_{j=1}^{t_{\hat{K}}-1} \phi_j(b_j) \cdot x_j^* + \phi_{t_{\hat{K}}}(b_{t_{\hat{K}}}) \cdot (\langle \vec{1}, \hat{x} \rangle - \langle \vec{1}, x^* \rangle) + \sum_{j=t_{\hat{K}}}^t x_j^* \\ &\geq \sum_{j=1}^{t_{\hat{K}}-1} \phi_j(b_j) \cdot x_j^* + \phi_{t_{\hat{K}}}(b_{t_{\hat{K}}}) \cdot (\sum_{j=t_{\hat{K}}}^t x_j^*) \\ &\geq \sum_{j=1}^t \phi_j(b_j) \cdot x_j^* = \langle \phi, x^* \rangle \end{aligned}$$

Note that the second inequality holds because $\langle \vec{1}, \hat{x} \rangle \geq \langle \vec{1}, x^* \rangle$, and the third inequality holds because bidders are numbered in the descending order of their virtual values. \square

LEMMA 5. For \mathcal{M}_2 , we have

$$\begin{aligned} &\sum_{k=1}^{\hat{K}} [\sum_{j=t_{k-1}+1}^{t_k-1} \phi_j(b_j) \cdot u_j + \phi_{t_k}(b_{t_k}) \cdot \min(u_{t_k}, N_k)] \\ &\geq \frac{1}{2} \{ \sum_{k=1}^{\hat{K}} [\sum_{j=t_{k-1}+1}^{t_k-1} \phi_j(b_j) \cdot u_j + \phi_{t_k}(b_{t_k}) \cdot N_k] \}. \end{aligned}$$

PROOF. For ease of references, we define the following notations for each $k \leq \hat{K}$,

$$\begin{aligned} \alpha_k &= \phi_{t_k}(b_{t_k}) \cdot \min(u_{t_k}, N_k) \\ \beta_k &= \sum_{j=t_{k-1}+1}^{t_k-1} \phi_j(b_j) \cdot u_j \\ \eta_k &= \phi_{t_k}(b_{t_k}) \cdot N_k \end{aligned}$$

Since $\sum_{j=t_{k-1}+1}^{t_k-1} u_j + \min(u_{t_k}, N_k) \geq N_k$ and for all $j \in \{t_{k-1} + 1, \dots, t_k\}$, $\phi_j(b_j) \geq \phi_{t_k}(b_{t_k})$, we have $\alpha_k + \beta_k \geq \eta_k$. It is clear $\alpha_k + \beta_k \geq \alpha_k$ due to the non-negativity of β_k .

Thus we have $(\alpha_k + \beta_k) \geq \frac{1}{2}(\alpha_k + \eta_k)$. By summing over k , we have the lemma proved. \square

THEOREM 6. When the value of bidder i follows a known regular distribution F_i , and $\forall i, l_i \leq N_K$, \mathcal{M}_2 achieves a competitive ration of at least $\frac{1}{4}$.

PROOF. The expected allocation profile of \mathcal{M}_2 is

$$\hat{x}' = (\frac{u_1}{2}, \dots, \frac{\min(u_{t_1}, N_1)}{2}, \dots, \frac{\min(u_{t_{\hat{K}}}, N_{\hat{K}})}{2}, 0, \dots, 0).$$

According to Lemma 4 and 5, we have that the virtual surplus $\langle \phi, \hat{x}' \rangle \geq \frac{1}{4}OPT(\mathbf{N}, b)$. Then the theorem is proven by taking expectation over the valuation/bid profile. \square

REMARK 7. The assumption, $l_i \leq N_K, \forall i$, means that the minimum demand of a merchant can be satisfied by displaying her deal to all the web users at the K -th slot. This assumption is reasonable since common merchants will not set the minimum demands too high so as not be rejected by the publisher.

5. PRIOR-INDEPENDENT SETTING

In this section, we focus on the prior-independent setting, in which the valuations of all the bidders are i.i.d. drawn from an unknown distribution F . We assume that the number of bidders is sufficiently large so that the empirical distribution F' estimated from their bids (given that they make truthful bidding) can approximate the real distribution F very well.

5.1 Two Mechanisms

Based on \mathcal{M}_1 , we can easily obtain a new mechanism \mathcal{M}_3 with a competitive ratio of $\frac{\beta-1}{\beta}$ (here $\beta > 1$ is an input parameter) for the prior-independent setting.

The new mechanism \mathcal{M}_3 consists of three steps. (1) Partition bids b at random into two sets: for each bid, with probability $\frac{1}{\beta}$ put it in set b' and otherwise b'' . (2) Calculate the empirical distribution F' using the bid set b' . (3) Run \mathcal{M}_1 with all the effective impressions, distribution F' and bid set b'' .

The estimation of the empirical distribution F' in \mathcal{M}_3 is independent of b'' , the bids of the bidders included in the auction. Therefore, given \mathcal{M}_1 is Bayesian incentive compatible, \mathcal{M}_3 is also Bayesian incentive compatible. Furthermore, it is not difficult to prove Theorem 8 about the competitive ratio of the mechanism. We omit the proof here because of space limitations.

THEOREM 8. If the valuations of the bidders are independently drawn from a regular distribution F and $\frac{M}{\beta}$ is statistically large enough, \mathcal{M}_3 achieves the competitive ratio of $\frac{\beta-1}{\beta}$.

Similar to the case in the Bayesian setting, although \mathcal{M}_3 has a good competitive ratio, it suffers from the computational inefficiency in practice. Here we design a simple greedy mechanism \mathcal{M}_4 which can achieve a much lower computational complexity, by leveraging \mathcal{M}_2 as a building block.

\mathcal{M}_4 is very similar to \mathcal{M}_3 . It also consists of three steps. The only difference between \mathcal{M}_3 and \mathcal{M}_4 is that we run \mathcal{M}_2 instead of \mathcal{M}_1 in the third step of \mathcal{M}_4 .

It is easy to verify that \mathcal{M}_4 is also Bayesian incentive compatible, and to obtain the following theorem about the competitive ratio of \mathcal{M}_4 . The detailed proof is omitted due to its simplicity.

THEOREM 9. If the valuations of the bidders are independently drawn from a regular distribution F , $\frac{M}{\beta}$ is statistically large enough, and $\forall i, l_i \leq N_K$, \mathcal{M}_4 achieves a competitive ratio $\frac{\beta-1}{4\beta}$.

5.2 Robust Approximation Mechanisms

In this subsection, we discuss a seemingly unobvious but practically important issue regarding mechanisms \mathcal{M}_3 and \mathcal{M}_4 . That is, both mechanisms rely on random partition of the bidders, and thus might not be very robust in the following sense: although their competitive ratios are guaranteed in expectation, the revenue of the worst-case instantiation of the mechanism can be arbitrarily low as compared with the optimal revenue. For example, suppose $\beta = 2$, $K = 1$, $N_1 = 1000000$, $M = 200$, and every bidder's value is 1 with probability 1. Let us consider a specific partition b' and b'' : $l_i = 10000, u_i = 20000$ for the 100 bidders in b' and $l_i = u_i = 1$ for the 100 bidders in b'' . For such a partition, the revenue achieved by \mathcal{M}_3 is 100 while the Bayesian optimal revenue is 1000000, which corresponds to a very low competitive ratio.

To be formal, we introduce a concept called *robust ratio* to evaluate the worst-case performance of a mechanism, as defined below.

DEFINITION 10. Let $OPT(F)$ denotes the optimal expected revenue of an incentive compatible mechanism under distribution F . Let σ be a random instance of the mechanism \mathcal{A} , and $Rev_\sigma(\mathcal{M}, F)$ be the expected revenue³ of the random instance σ of \mathcal{M} under distribution F . Then the robust ratio of a mechanism \mathcal{M} is defined as $\inf_{\sigma, F} \frac{Rev_\sigma(\mathcal{M}, F)}{OPT(F)}$.

Note that the robust ratio ranges in $[0, 1]$, and the robust ratio of a mechanism is no larger than its competitive ratio.

To make mechanisms \mathcal{M}_3 and \mathcal{M}_4 robust (i.e., with a bounded robust ratio), we propose a technique called *supply partition*, i.e., partitioning the impressions. Due to space restrictions, we just take the modification to \mathcal{M}_3 as an example and present a new mechanism \mathcal{M}_5 as below.

Algorithm 3 A Robust Approximation Mechanism (\mathcal{M}_5)

- 1: Partition bids b at random into two sub sets: for each bid, with probability $\frac{1}{2}$ put it in b' and otherwise b'' .
 - 2: Calculate an empirical distribution F' using the bid set b' and an empirical distribution F'' using the bid set b'' .
 - 3: **for** each slot i **do**
 - 4: If i is odd, offering the first $N/2$ impressions to set b' and the second $N/2$ impressions to set b'' .
 - 5: If i is even, offering the first $N/2$ impressions to set b'' and the second $N/2$ impressions to set b' .
 - 6: **end for**
 - 7: Construct a new supply vector $\mathbf{N}' = (\frac{N_1+N_2}{2}, \frac{N_3+N_4}{2}, \dots, \frac{N_{\lceil K/2 \rceil * 2 - 1} + N_{\lceil K/2 \rceil * 2}}{2}, 0, \dots, 0)$.
 - 8: Run \mathcal{M}_1 with the new supply vector \mathbf{N}' , distribution F'' and bid set b' . Run \mathcal{M}_1 with the new supply vector \mathbf{N}' , distribution F' and bid set b'' .
-

Note in \mathcal{M}_5 , though the number of impressions allocated to bidders in slot i is $N/2$, the actual effective impressions is $N_i/2$ because of different conversion rate for each slot. It is easy to see that \mathcal{M}_5 is Bayesian incentive compatible. The following theorem shows that it has guarantee on both competitive ratio and robust ratio.

³Note that for “expected revenue”, the expectation is taken over the randomness of bidders' valuations but not over the random instances of the mechanism.

THEOREM 11. If bidders' valuations are independently drawn from a regular distribution F , the number of bidders is statistically large enough, and $l_i \leq \frac{N_1+N_2}{2}, \forall i$, \mathcal{M}_5 achieves a competitive ratio of $\frac{3}{7}$ and a robust ratio of $\frac{1}{4}$.

PROOF. To ease the proof, we give some notations first. Given the value profile v and the majorization constraint vector $\mathbf{N} = (N_1, N_2, \dots, N_K)$, we denote (x_1^*, \dots, x_M^*) as the optimal allocation vector and $OPT(\mathbf{N}, b) = \sum_{i \in [M]} \phi_i \cdot x_i^*$ as the optimal virtual surplus⁴ generated from the bidder set b when the constraint vector \mathbf{N} is imposed and when F is known.

First, we prove the robust ratio of the mechanism.

For any partition (A, B) of the M bidders, we denote $OPT_A(\mathbf{N}, b) = \sum_{i \in A} \phi_i \cdot x_i^*$ and $OPT_B(\mathbf{N}, b) = \sum_{i \in B} \phi_i \cdot x_i^*$.

Let $OPT(\mathbf{N}', C)$ be the optimal virtual surplus generated from the bidder set C and the majorization constraint vector \mathbf{N}' , where C can be either A or B . As long as we can construct a feasible allocation under \mathbf{N}' and C , whose virtual surplus is a constant approximation to $OPT_C(\mathbf{N}, b)$, we can say that $OPT(\mathbf{N}', C)$ is also a constant approximation to $OPT_C(\mathbf{N}, b)$.

We re-number the bidders in C according to the descending order of their allocated numbers of impressions: $x_1^* \geq x_2^* \geq \dots \geq x_{|C|}^*$, where $|C|$ is the number of bidders in C . We focus on bidder 1 and consider the following two cases.

Case 1: If $\phi_1 x_1^* \geq \frac{N_1}{2 \cdot N_1 + N_2} \cdot OPT_C(\mathbf{N}, b)$, we allocate $\min(x_1^*, \frac{N_1+N_2}{2})$ effective impressions to bidder 1 under the constraint \mathbf{N}' , and the virtual surplus of this allocation is:

$$\begin{aligned} \phi_1 \cdot \min(x_1^*, \frac{N_1 + N_2}{2}) &= \frac{\min(x_1^*, \frac{N_1+N_2}{2})}{x_1^*} \cdot \phi_1 \cdot x_1^* \\ &\geq \frac{N_1 + N_2}{2N_1} \cdot \frac{N_1}{2 \cdot N_1 + N_2} \cdot OPT_C(\mathbf{N}, b) \\ &= \frac{1}{2} \cdot \frac{N_1 + N_2}{2 \cdot N_1 + N_2} \cdot OPT_C(\mathbf{N}, b) \end{aligned}$$

Case 2: If $\phi_1 x_1^* < \frac{N_1}{2 \cdot N_1 + N_2} \cdot OPT_C(\mathbf{N}, b)$, under \mathbf{N}' , for each pair of bidders $(2i, 2i+1)$, $i \in \{1, \dots, \lfloor |C|/2 \rfloor\}$, if $\phi_{2i} \cdot x_{2i}^* \geq \phi_{2i+1} \cdot x_{2i+1}^*$, we allocate x_{2i}^* effective impressions to bidder $2i$; otherwise, we allocate x_{2i+1}^* effective impressions to bidder $2i+1$. Since $\{x_i^*\}$ are sorted and satisfy the majorization constraint vector \mathbf{N} , it is easy to see that for all $i \in \{1, \dots, \lfloor K/2 \rfloor\}$,

$$\begin{aligned} \sum_{j=1}^i \max(x_{2j}^*, x_{2j+1}^*) &= \sum_{j=1}^i x_{2j}^* \leq \sum_{j=1}^i \frac{N_{2j-1} + N_{2j}}{2}, \\ \sum_{j=1}^{\lfloor |C|/2 \rfloor} \max(x_{2j}^*, x_{2j+1}^*) &\leq \sum_{j=1}^{\lfloor |C|/2 \rfloor} x_{2j}^* \leq \sum_{j=1}^{\lfloor K/2 \rfloor} \frac{N_{2j-1} + N_{2j}}{2}. \end{aligned}$$

Therefore the allocation is feasible under \mathbf{N}' . Furthermore, we can lower bound the virtual surplus VS of this allocation:

$$VS \geq \frac{1}{2} \sum_{j=2}^{|C|} \phi_j x_j^* \geq \frac{1}{2} \cdot \frac{N_1 + N_2}{2 \cdot N_1 + N_2} \cdot OPT_C(\mathbf{N}, b).$$

By jointly considering the two cases, for any partition (A, B) of the bidders, we can bound the total virtual surplus

⁴We can run the Bayesian optimal auction (i.e. \mathcal{M}_1) to get the optimal allocation vector and the optimal virtual surplus.

of \mathcal{M}_5 as follows,

$$\begin{aligned} Rev_{(A,B)}(\mathcal{M}_5) &\geq OPT(\mathbf{N}', A) + OPT(\mathbf{N}', B) \\ &\geq \frac{1}{2} \cdot \frac{N_1 + N_2}{2 \cdot N_1 + N_2} \cdot (OPT_A(\mathbf{N}, b) + OPT_B(\mathbf{N}, b)) \\ &= \frac{1}{2} \cdot \frac{N_1 + N_2}{2 \cdot N_1 + N_2} \cdot OPT(\mathbf{N}, b) \geq \frac{1}{4} \cdot OPT(\mathbf{N}, b). \end{aligned}$$

By taking expectation of the leftmost side and the rightmost side of the above inequality over the distribution of bidders' valuations, we complete the proof of the robust ratio.

Second, we prove the competitive ratio of the mechanism.

For this purpose, we re-number the bidders as $x_1^* \geq x_2^* \geq \dots \geq x_M^*$. Denote $\alpha = \frac{\phi_1 x_1^*}{OPT(\mathbf{N}, b)}$. Recall that $OPT(\mathbf{N}', A) + OPT(\mathbf{N}', B)$, the virtual surplus generated by the mechanism, is the optimal virtual surplus under partition (A, B) and the majorization constraint \mathbf{N}' . Given \mathbf{N}' and (A, B) , the virtual surplus of any feasible allocation lower bounds the virtual surplus of the mechanism. The basic idea of our proof is that we construct two feasible allocations, both of which can lower bound the virtual surplus of the mechanism, and show that the maximum of the two lower bounds is a constant approximation of $OPT(\mathbf{N}, b)$.

Allocation 1: No matter bidder 1 is partitioned to set A or B , we allocate $\min(x_1^*, \frac{N_1 + N_2}{2})$ effective impressions to him/her. We then pair every adjacent two remaining bidders: $(2i, 2i + 1)$, $i \in \{1, \dots, \lfloor M/2 \rfloor\}$ and make allocation for each pair depending on which set the two bidders are partitioned into:

- (1) With probability $1/4$, bidder 1 and $2i$ are partitioned in the same set and bidder $2i + 1$ is in the other set. For this case, we allocate x_{2i+1}^* effective impressions to bidder $2i + 1$.
- (2) With probability $1/4$, bidder 1 and $2i + 1$ are in the same set and bidder $2i$ is in the other set. For this case, we allocate x_{2i}^* effective impressions to bidder $2i$.
- (3) With probability $1/4$, bidder $2i$ and $2i + 1$ are in the same set and bidder 1 is in the other set. For this case, if $\phi_{2i} \cdot x_{2i}^* > \phi_{2i+1} \cdot x_{2i+1}^*$, we allocate x_{2i}^* effective impressions to bidder $2i$; otherwise, we allocate x_{2i+1}^* effective impressions to bidder $2i + 1$.
- (4) With probability $1/4$, bidder 1, $2i$ and $2i + 1$ are partitioned into the same set. For this case, we do not allocate any impression to bidders $2i$ and $2i + 1$.

In expectation, the virtual surplus (denoted as $VS(2i, 2i + 1)$) generated from the pair $(2i, 2i + 1)$ can be lower bounded as below.

$$\begin{aligned} VS(2i, 2i + 1) &\geq \frac{1}{4} \cdot \{\phi_{2i} \cdot x_{2i}^* + \phi_{2i+1} \cdot x_{2i+1}^*\} \\ &+ \frac{1}{2} (\phi_{2i} \cdot x_{2i}^* + \phi_{2i+1} \cdot x_{2i+1}^*) = \frac{3}{8} \cdot (\phi_{2i} \cdot x_{2i}^* + \phi_{2i+1} \cdot x_{2i+1}^*) \end{aligned}$$

For the set in which bidder 1 is partitioned to, since we only allocate impressions to bidder 1 and $l_1 \leq \frac{N_1 + N_2}{2}$, the allocation is definitely feasible. In the other set, since for all i :

$$\begin{aligned} \sum_{j=1}^i x_{2j}^* &\leq \frac{1}{2} \sum_{j=1}^{2i} x_j^* \leq \frac{1}{2} \sum_{j=1}^{2i} N_j, \\ \sum_{j=1}^{\lfloor M/2 \rfloor} x_{2j}^* &\leq \frac{1}{2} \sum_{j=1}^M x_j^* \leq \frac{1}{2} \sum_{j=1}^K N_j, \end{aligned}$$

the allocation is also feasible with respect to \mathbf{N}' . The virtual surplus VS_1 of this allocation can be lower bounded as:

$$\begin{aligned} VS_1 &= [\min(x_1^*, \frac{N_1 + N_2}{2})\phi_1 + \sum_{i=1}^{\lfloor M/2 \rfloor} VS(2i, 2i + 1)] \\ &\geq [\frac{\min(x_1^*, \frac{N_1 + N_2}{2})}{x_1^*} \alpha + \frac{3}{8}(1 - \alpha)] \cdot \Phi^* \\ &\geq [\frac{N_1 + N_2}{2N_1} \alpha + \frac{3}{8}(1 - \alpha)] \cdot OPT(\mathbf{N}, b). \end{aligned}$$

Allocation 2: We ignore bidder 1 and consider each pair of bidders $(2i, 2i + 1)$, $i \in \{1, \dots, \lfloor M/2 \rfloor\}$.

- (1) With probability $1/2$, two bidders in the same pair are located in the same set. If $\phi_{2i} \cdot x_{2i}^* \geq \phi_{2i+1} \cdot x_{2i+1}^*$, we allocate x_{2i}^* effective impressions to bidder $2i$; otherwise, we allocate x_{2i+1}^* effective impressions to bidder $2i + 1$.
- (2) With probability $1/2$, two bidders in the same pair are in different sets. We allocate x_{2i}^* effective impressions to bidder $2i$ and x_{2i+1}^* effective impressions to bidder $2i + 1$.

The above allocation still satisfies the majorization constraint \mathbf{N}' . The virtual surplus VS_2 of this allocation can be lower bounded as below.

$$\begin{aligned} VS_2 &\geq \frac{1}{2} \cdot (1 - \alpha) \cdot OPT(\mathbf{N}, b) + \frac{1}{2} \cdot \frac{1}{2} \cdot (1 - \alpha) \cdot OPT(\mathbf{N}, b) \\ &= \frac{3}{4} \cdot (1 - \alpha) \cdot OPT(\mathbf{N}, b) \end{aligned}$$

Denote $h(\alpha) = \max\{\frac{N_1 + N_2}{2N_1} \alpha + \frac{3}{8}(1 - \alpha), \frac{3}{4} \cdot (1 - \alpha)\}$. It is easy to see that $\frac{3}{4} \cdot (1 - \alpha)$ is a decreasing function of α . Since $\frac{N_1 + N_2}{2N_1} > \frac{3}{8}$, $\frac{N_1 + N_2}{2N_1} \alpha + \frac{3}{8}(1 - \alpha)$ is an increasing function of α . Therefore $h(\alpha)$ achieves its minimal value when the two terms are equal, i.e., $\frac{N_1 + N_2}{2N_1} \alpha + \frac{3}{8}(1 - \alpha) = \frac{3}{4} \cdot (1 - \alpha)$. Thus, we get

$$h(\alpha) \geq \frac{3(N_1 + N_2)}{4(N_1 + N_2) + 3N_1} = \frac{3}{7}.$$

That is, for any given bid/valuation profile b , the revenue of the mechanism averaged over the random partitions is lower bounded by $\frac{3}{7}OPT(\mathbf{N}, b)$. By further taking expectation over the value profile, we complete the proof of the competitive ratio. \square

6. FUTURE WORK

As for future work, we plan to investigate the following aspects. First, we may prove the upper bound of competitive ratio for any polynomial mechanism in our model, thus to show our mechanisms are nearly the best that can do, or we may build other elegant mechanisms with better competitive ratio. Second, our approach for approximation mechanism may have significance on its own, we may find other interesting applications for our general analysis scheme. Third, we have assumed that the number of bidders is statistically large enough to get an accurate estimation of the distribution of bidders' valuations for the prior-independent setting. We will quantify how the number of bidders affect the revenue of the designed mechanisms. Finally, we have only considered offline auctions in this work, where the effective impressions in a future time period is known when the auction is executed. A more practical setting is online auctions.

REFERENCES

- [1] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 1–7. ACM, 2006.
- [2] G. Aggarwal and J. D. Hartline. Knapsack auctions. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1083–1092. ACM, 2006.
- [3] N. Alon, F. Fischer, A. D. Procaccia, and M. Tennenholtz. Sum of us: Truthful self-selection. In *TARK*.
- [4] A. Archer, C. Papadimitriou, K. Talwar, and É. Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics*, 1(2):129–150, 2004.
- [5] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 482–491. IEEE, 2001.
- [6] J. Byers, M. Mitzenmacher, and G. Zervas. Daily deals: Prediction, social diffusion, and reputational ramifications. In *Proceedings of the fifth ACM international conference on Web search and data mining*, pages 543–552. ACM, 2012.
- [7] J. Byers, M. Mitzenmacher, and G. Zervas. The groupon effect on yelp ratings: a root cause analysis. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 248–265. ACM, 2012.
- [8] Y. Cai, M. Mahdian, A. Mehta, and B. Waggoner. Designing markets for daily deals. In *Web and Internet Economics*, pages 82–95. Springer, 2013.
- [9] G. B. Dantzig. Discrete-variable extremum problems. *Operations Research*, 5(2):266–288, 1957.
- [10] B. C. Dean, M. X. Goemans, and J. Vondrck. Approximating the stochastic knapsack problem: The benefit of adaptivity. In *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on*, pages 208–217. IEEE, 2004.
- [11] U. Dholakia. What makes groupon promotions profitable for businesses? Available at SSRN 1790414, 2011.
- [12] B. Edelman, S. Jaffe, and S. Kominers. To groupon or not to groupon: The profitability of deep discounts. *Harvard Business School NOM Unit Working Paper*, (11-063), 2011.
- [13] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Technical report, National Bureau of Economic Research, 2005.
- [14] A. V. Goldberg, J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [15] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 735–744. Society for Industrial and Applied Mathematics, 2001.
- [16] R. Gonen and D. Lehmann. Optimal solutions for multi-unit combinatorial auctions: Branch and bound heuristics. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 13–20. ACM, 2000.
- [17] T. Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.
- [18] J. D. Hartline. Approximation in economic design. *Lecture Notes*, 2012.
- [19] W. Kong, J. Li, T.-Y. Liu, and T. Qin. Optimal allocation for chunked-reward advertising. In *Web and Internet Economics*, pages 291–304. Springer, 2013.
- [20] W. Kong, J. Li, T. Qin, and T.-Y. Liu. Revenue optimization for group-buying websites. *arXiv preprint arXiv:1305.5946*, 2013.
- [21] Y. Lvovskaya, S. Tan, and C. Zhong. *Online Discount Coupon Promotions & Repurchasing Behaviors: The Groupon Case*. PhD thesis, Mälardalen University, 2012.
- [22] E. Maskin, J. Riley, and F. Hahn. Optimal multi-unit auctions. *The economics of missing markets, information, and games*, 1989.
- [23] P. R. Milgrom and R. J. Weber. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122, 1982.
- [24] R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [25] T. Roughgarden. Cs364a: Algorithmic game theory lecture# 4: Algorithmic mechanism design. 2013.
- [26] M. Song, E. Park, B. Yoo, and S. Jeon. Is the daily deal social shopping?: An empirical analysis of purchase and redemption time of daily-deal coupons. *Working paper*, 2012.
- [27] D. Zhao, W. Chen, and Q. Ye. Consumer purchase behavior under price promotion: Evidence from groupon daily deals. In *SIGBPS Workshop on Business Processes and Services (BPS’12)*, page 141, 2012.