

Mechanism Design for Crowdsourcing: An Optimal $1-1/e$ Competitive Budget-Feasible Mechanism for Large Markets

Nima Anari*, Gagan Goel†, Afshin Nikzad‡

Abstract

In this paper we consider a mechanism design problem in the context of large-scale crowdsourcing markets such as Amazon’s Mechanical Turk (MTRK), ClickWorker (CLKWRKR), CrowdFlower (CRDFLWR). In these markets, there is a requester who wants to hire workers to accomplish some tasks. Each worker is assumed to give some utility to the requester on getting hired. Moreover each worker has a minimum cost that he wants to get paid for getting hired. This minimum cost is assumed to be private information of the workers. The question then is - if the requester has a limited budget, how to design a direct revelation mechanism that picks the right set of workers to hire in order to maximize the requester’s utility?

We note that although the previous work (Singer (2010); Chen et al. (2011)) has studied this problem, a crucial difference in which we deviate from earlier work is the notion of *large-scale* markets that we introduce in our model. The notion of a large-scale market that we consider is a natural one which states that the (private) cost of each worker is *small* compared to the budget of the requester. Without the large market assumption, it is known that no mechanism can achieve a competitive ratio better than 0.414 and 0.5 for deterministic and randomized mechanisms respectively (while the best known deterministic and randomized mechanisms achieve an approximation ratio of 0.292 and 0.33 respectively). In this paper, we design a budget-feasible mechanism for large markets that achieves a competitive ratio of $1 - 1/e \simeq 0.63$. Our mechanism can be seen as a generalization of an alternate way to look at the *proportional share* mechanism, which is used in all the previous works so far on this problem. Interestingly, we can also show that our mechanism is optimal by showing that no truthful mechanism can achieve a factor better than $1 - 1/e$; thus, fully resolving this setting. Finally we consider the more general case of submodular utility functions and give new and improved mechanisms for the case when the market is large.

*University of California, Berkeley. anari@eecs.berkeley.edu

†Google Research, New York. gagangoel@google.com

‡Stanford University, Stanford. nikzad@stanford.edu

Contents

1	Introduction	1
1.1	The Model	2
1.1.1	The Large Market Assumption	3
1.2	Our Results	4
1.3	Related Work	5
1.4	Roadmap	5
2	A Simple $\frac{1}{2}$-Approximate Truthful Mechanism	6
3	Our Approach	7
3.1	A Notion of An Allocation Rule	7
3.2	Payment Rule	8
3.3	Example	9
3.4	First Attempt: A Parameterized Class of Envy-Free Mechanisms	10
4	A Parameterized Class of Truthful Mechanisms	11
5	A $(1 - 1/e)$-Approximate Optimal Truthful Mechanism	12
5.1	Finding an optimal f for the (non-truthful) Mechanism Envy-Free(f)	13
5.1.1	Preliminaries	13
5.1.2	Approximation Ratio of Mechanism Envy-Free(f)	14
5.2	Special case: Analyzing our truthful mechanism for unit utilities	15
6	Impossibility Result: On why $1 - 1/e$ is the best approximation possible	17
7	Conclusion	21
A	Analyzing Our Optimal Truthful Mechanism for General Additive Utilities (Continuation of Section 5)	24
B	Missing Proofs from Section 5.2	26
C	Mechanisms for Indivisible Items	27
C.1	Properties of the Rounding Procedure	27
C.2	Description of the Rounding Procedure	28
D	The Optimal Standard Allocation Rule	30
E	Submodular Utility Functions	33
E.1	Preliminaries	34
E.2	Basic Definitions	35
E.2.1	The Large Market Assumption	35
E.2.2	Oneway-Truthfulness	35
E.2.3	Submodular Functions	35
E.3	The Exp-Time Mechanism	36
E.4	An Improved Polynomial-Time Mechanism	40

E.5 From Oneway-Truthfulness to Truthfulness	44
E.6 Strictly Budget Feasible Mechanisms	47
F Hoeffding Bounds	49

1 Introduction

Crowdsourcing is a recent phenomenon that is used to describe the procurement of a large number of workers to do certain tasks. These tasks can be of a variety of natures, and - to give a few examples - include image annotation, data labeling for machine learning systems, consumer surveys, rating search engine results, spam detection, product reviews, etc. There are several platforms (such as Amazon’s Mechanical Turk (MTRK)) that facilitate and automate various steps involved in setting up and executing crowdsourcing tasks.

A key challenge in these online labor markets is to be able to properly price the tasks. Since the requester (the one who wants to procure workers) is usually budget constrained, pricing the tasks too high can result in lower output for the requester. On the other hand, pricing the tasks too low can disincentivize workers to work on the tasks. This makes pricing a non-trivial step for the requester when setting up a crowdsourcing task. One idea - to make pricing more automated and to prevent economic loss from poor pricing - is to design a direct revelation mechanism that solicits bids from workers to report their cost of participation, and based on this decide which workers to hire and how much to pay them.

A simple model that captures the above problem is as follows: There is a set S of workers. Worker i has a private cost c_i and provides utility u_i to the requester on getting hired. We want to design a truthful mechanism that decides which workers to recruit and how much to pay them. The goal is to maximize the requester’s utility without violating her budget constraint.

For the above model, Singer (2010) gave an incentive-compatible mechanism that achieves an approximation ratio¹ of $1/6$ compared to the *offline optimum* that knows the costs of the workers. Later on Chen et al. (2011) improved the approximation ratio to $\frac{1}{2+\sqrt{2}} \simeq 0.292$ (and to $1/3$ for randomized mechanisms). Chen et. al. also showed that no deterministic mechanism can achieve an approximation ratio better than $\frac{1}{1+\sqrt{2}} \simeq 0.414$, and no randomized mechanism can achieve an approximation ratio better than 0.5 .

Our work is motivated by the following observation: Most of the crowdsourcing tasks are *large-scale* in nature in terms of the number of workers involved. On the other hand if one looks at the impossibility result of Chen et al. (2011), they involve only a small number of workers (specifically, only 3 workers). Thus, this leads to a natural open question - *Do these lower bounds extend to the case of large markets? or can one design better mechanisms for this important case of large markets?*

In this paper, we seek to understand the above question. We show that one can significantly improve the approximation ratio for the case of large markets. We give a mechanism that achieves an approximation ratio of $1 - 1/e \simeq 0.63$ for large markets. In addition, we show that our mechanism is the best possible mechanism by showing that no truthful budget-feasible mechanism can achieve a

¹we use the terms approximation factor and competitive ratio interchangeably

factor better than $1 - 1/e$. Finally, we look at the more general case of submodular utility functions.

1.1 The Model

We define the model abstractly: Consider a reverse auction scenario with one buyer and n sellers, where the set of sellers is denoted by S . Each seller $i \in S$ owns a single item (denoted by item i) and has a *private* cost c_i for it. The buyer derives a utility of u_i from item i . The buyer has a limited budget B , and its goal is to buy a subset of items that maximizes her total utility ² without exceeding her budget.

Note that if the sellers are not strategic and the costs are known to the buyer, then this is the well-known *knapsack* optimization problem. However, the cost c_i is assumed to be a private information of seller i . Thus we are interested in designing direct-revelation mechanisms where the buyer solicits bids from the sellers, and then computes which sellers to buy from and how much to pay them. More formally, a mechanism \mathcal{M} consists of two functions $A : (\mathbb{R}_+)^n \rightarrow \{0, 1\}^n$ and $P : (\mathbb{R}_+)^n \rightarrow (\mathbb{R}_+)^n$. The allocation function $A(\cdot)$ takes as input the costs of n sellers and reports the set of winners. The payment function $P(\cdot)$ takes as input the costs of n sellers and reports how much is paid to each seller. We will use functions $A_i : (\mathbb{R}_+)^n \rightarrow \{0, 1\}$ and $P_i : (\mathbb{R}_+)^n \rightarrow \mathbb{R}_+$, for each $i \in S$, to refer to the restriction of functions $A(\cdot)$ and $P(\cdot)$ to seller i . In other words, A_i, P_i represent the i -th element of the output of functions A, P , respectively.

The mechanism $\mathcal{M} = (A, P)$ should satisfy the following properties:

1. Budget Feasibility: The sum of the payments made to the sellers should not exceed B , i.e., $\sum_i P_i(\mathbf{c}) \leq B$ for all \mathbf{c} .
2. Individual rationality: A winner $i \in S$ is paid at least c_i .
3. Truthfulness/Incentive-Compatibility: Reporting the true cost should be a dominant strategy of the sellers, i.e. for all non-truthful reports \bar{c}_i from seller i , it holds that

$$P_i(\bar{c}_i, \mathbf{c}_{-i}) - c_i \cdot A_i(\bar{c}_i, \mathbf{c}_{-i}) \leq P_i(c_i, \mathbf{c}_{-i}) - c_i \cdot A_i(c_i, \mathbf{c}_{-i})$$

Among all mechanisms that satisfy the above properties, we are interested in the ones that give high utility to the buyer. Note that no mechanism can achieve utility larger than $U^*(\mathbf{c}, \mathbf{u})$, where $U^*(\mathbf{c}, \mathbf{u})$ is the utility of the knapsack optimization problem assuming costs of the sellers are known to the buyer. When there is no risk of confusion, we also denote $U^*(\mathbf{c}, \mathbf{u})$ by U^* for brevity.

We say a mechanism \mathcal{M} is an α -approximation (for $\alpha \leq 1$) if it gives utility at least $\alpha \cdot U^*(\mathbf{c}, \mathbf{u})$ for any \mathbf{c} and \mathbf{u} .

²We note that total utility is only a function of the items that the buyer buys, and not of the amount she spends on it.

Indivisible vs Divisible Items. Note that the above description is given for indivisible items, however, we can define the above problem for divisible items as well. For instance, if the item being sold by a seller is his own time, then it can be modeled as a divisible item. For fraction $x \leq 1$ of a divisible item, the cost of seller i is $x \cdot c_i$ and the utility obtained by the buyer is $x \cdot u_i$. The allocation function for divisible items is defined as $A : (\mathbb{R}_+)^n \rightarrow [0, 1]^n$.

More general utility functions. An interesting generalization of the above model is when the utility function over the set of items is a submodular function rather than an additive function. We denote this function by $U : 2^S \rightarrow \mathbb{R}_+$ (for additive functions, $U(T) = \sum_{i \in T} u_i$, for $\forall T \subseteq S$). We assume that the utility function is known to the buyer.

1.1.1 The Large Market Assumption

Crowd-sourcing systems are excellent examples of *large markets*. Informally speaking, a market is said to be large if the number of participants are large enough that no single person can affect the market outcome significantly. Our results take advantage of this nature of crowdsourcing markets to give better mechanisms.

We define the **large market assumption** as follows: We assume that in our model, the cost of a single item is very small compared to the buyer's budget B . More formally, let $c_{\max} = \max_{i \in S} \{c_i\}$. Then, the large market assumption is defined as below.

The Large Market Assumption: $c_{\max} \ll B$.

In other words, we define the *largeness ratio* of the market to be $\theta = \frac{c_{\max}}{B}$ and analyze our mechanisms for when $\theta \rightarrow 0$.

This assumption - also known as the *small bid to budget ratio assumption* - is used in other large-market problems as well (for instance, see Mehta et al. (2007) for a similar definition with application in online advertising). All the mechanisms that we present in the main body of the paper (mechanisms for additive utility functions) will be analyzed under this assumption. The mechanisms that we design for submodular utility functions work under a different large market assumption which is explained below.

An Alternative Assumption We also suggest another definition for large markets, the discussion of which will be deferred to the appendix. Our mechanisms for submodular utility functions work under this assumption; moreover, we can slightly modify our mechanisms for additive utility functions so that they work under this assumption as well, while preserving their approximation ratio. We define this assumption below.

Let $u_{\max} = \max_{i \in S} u_i$ and U^* be the utility of the optimum solution (i.e. the maximum utility that the buyer can achieve when the costs are known to her). This large market assumption states

that:

An Alternative Large Market Assumption: $u_{\max} \ll U^$.*

In other words, we define the largeness ratio of the market to be $\theta = \max_{i \in S} \frac{u_i}{U^*}$ and analyze our mechanisms for when $\theta \rightarrow 0$.

We note that our impossibility result for additive utilities (Section 6) holds for either of the two definitions.

1.2 Our Results

In this paper, we design optimal budget-feasible mechanisms for *large markets*. To the best of our knowledge, we are the first ones to study the case of large markets. We list our results below:

1. If the items are divisible, we design a deterministic mechanism which satisfies all the required properties (i.e. budget feasibility, individual rationality and truthfulness) and has an approximation ratio of $1 - 1/e$ (Section 5). Note that previously, no mechanism was known for the case of divisible items. In fact, one can show that no bounded approximation ratio is possible for divisible items if the large market assumption is dismissed.
2. If the items are indivisible, we can modify our mechanism and give a randomized truthful mechanism for this case which achieves approximation ratio $1 - 1/e$. (Section C)
3. In Section 6, we show that the above results are optimal by proving that no truthful (and possibly) randomized mechanism can achieve approximation ratio better than $1 - 1/e$. Our impossibility result holds for both cases of divisible and indivisible items.
4. For the case of submodular utility functions, we design deterministic mechanisms that achieve approximation ratios of $\frac{1}{2}$ and $\frac{1}{3}$ with exponential and polynomial running times respectively. Note that we only consider the case of indivisible items for submodular utility functions. (Section E)

As we saw in Section 1.1.1, one could define a notion of θ -large market, i.e. a market with largeness ratio θ . To gain a better understating of the problem, we focus on large markets (i.e. when $\theta \rightarrow 0$) and state our main theorems for this setting. However, our mechanisms do not need “very large” markets to perform well; for instance, in the knapsack problem with additive utilities, the approximation ratio ³ is $(1 - 1/e) \cdot (1 - 6\theta/5)$ when all the items have equal utilities (Section 5.2). Thus, say for $\theta = 1/20$ and $\theta = 1/40$ (which are reasonable assumptions in many settings) we get approximation factors 0.592 and 0.613 respectively.

³we didn't try to optimize the dependence on θ in our analysis as we focus on the main ideas for the sake of better understanding.

Also we point out that the above results have applications beyond crowdsourcing - for instance, see Singer (2011) for application in marketing over social networks, and Horel et al. (2013) for application in experiment design. Singer (2011) provides a truthful mechanism with approximation ratio ≈ 0.032 and Horel et al. (2013) provides an approximately truthful mechanism with approximation ratio ≈ 0.077 . For both these settings, large market assumption is a very reasonable assumption to make; thus, our results apply to these applications as well. In particular, our results give fully truthful mechanisms for these applications with approximation ratios $\frac{1}{2}, \frac{1}{3}$ (for exponential and polynomial running time respectively) in large markets.

1.3 Related Work

The most relevant related work is that of Singer (2010) and Chen et al. (2011). Singer (2010) first introduced this model (without the large market assumption). For the case of additive utilities and indivisible items, he gave a deterministic mechanism with an approximation ratio of $1/6$. Chen et al. (2011) later improved it to $1/(2 + \sqrt{2})$, and also gave a randomized mechanism with an approximation ratio of $1/3$. Chen et al. (2011) and Singer (2010) gave lower bounds of $1/(1 + \sqrt{2})$ and $1/2$ for deterministic and randomized mechanisms respectively. For the case of submodular utilities, Singer (2010) gave a randomized mechanism with an approximation ratio of $1/112$ which was improved to $1/7.91$ by Chen et al. (2011). Chen et al. (2011) also gave an exponential time deterministic mechanism for submodular utility functions with an approximation ratio of $1/8.34$.

Dobzinski et al. (2011) looked at the more general sub-additive utility functions and gave a $1/\log^2(n)$ and $1/\log^3(n)$ approximation ratio for randomized and deterministic mechanisms respectively. Singla and Krause (2013a) design budget-feasible mechanisms for adaptive submodular functions with applications in community sensing.

In another work, Bei et al. (2012) study this problem in Bayesian setting. Singer (2011) looks at the application of this model in marketing over social networks. Horel et al. (2013) study the application of this model in experiment design.

Another related model that has been inspired from crowdsourcing applications is when the workers arrive online. A sequence of papers model this as an online learning problem. See Singla and Krause (2013b); Badanidiyuru et al. (2012); Singer and Mittal (2013) for more details.

Finally, we note that our assumption for large markets is similar to the assumption made in other application areas; notably in the Adwords problem as studied by Mehta et al. (2007). See Goel and Mehta (2008); Devanur and Hayes (2009); Feldman et al. (2010, 2009) for other models motivated by online advertising where they make similar assumptions.

1.4 Roadmap

The readers are encouraged to read this section before proceeding further. We begin by presenting a $\frac{1}{2}$ -approximate mechanism (in section 2) using a simple *proportional share mechanism* which was

proposed in Singer (2010); Chen et al. (2011). Then, in section 3, we build our new ideas that help us find the optimal mechanism for the case of the additive utilities. This section also defines a simple version of our mechanism that is not truthful but is envy-free. We define our truthful mechanism in section 4. Both our envy-free and truthful mechanisms are parameterized by a single input variable (which we call *standard allocation rule*). Most of the work in the later sections goes in finding (and proving) the optimal standard allocation rule. In particular, in section 5, we find the standard allocation rule for which our mechanism provides an approximation ratio of $1 - 1/e$ in large markets. In section 6, we complement this result by showing that no truthful mechanism can achieve approximation ratio better than $1 - 1/e$. In section C, we adapt our mechanism to the case of indivisible items. In section E, we present two mechanisms for submodular utility functions which have exponential and polynomial running times and approximation ratios $\frac{1}{2}$ and $\frac{1}{3}$, respectively.

2 A Simple $\frac{1}{2}$ -Approximate Truthful Mechanism

In this section, we briefly explain the previous mechanism designed for this problem for the additive utility functions that gives an approximation factor of $\frac{1}{2}$ in large markets.

Definition 1 Cost-per-utility rate of a seller i is equal to c_i/u_i .

A natural approach to this problem tries to find a single *payment-per-utility rate* (denoted by rate r) at which all the winning sellers get paid. In other words, this approach picks a single number r and makes a payment of $r \cdot u_i$ to seller i if she wins and pays her 0 otherwise. For brevity, we sometimes call the payment-per-utility rate r simply the *rate* r when there is no risk of confusion.

Individual rationality implies that a seller i is willing to sell her item at rate r iff $r \geq c_i/u_i$. Initially the buyer declares a very large rate r , and then sees which sellers are willing to sell at this rate. If the total cost to buy from all these sellers at rate r is higher than the budget B , then the buyer decreases the rate r . More formally, a natural descending price auction for this problem works as follows:

1. Let A denote the set of active sellers, and initially set $A = S$.
2. Start with a very high rate r .
3. Verify if all the active sellers can be paid with rate r , i.e. whether $\sum_{i \in A} r \cdot u_i \leq B$ or not.
4. If the payment is feasible, then allocate the subset A , make the payment and stop.
5. If the payment is not feasible then decrease r slightly; update A accordingly by removing the sellers i for whom $c_i/u_i > r$; go to Step 3.

The above auction captures the main idea behind the *proportional share* mechanisms designed in Singer (2010); Chen et al. (2011)⁴, although they describe it in a forward auction format. It is not hard to see that the above mechanism is truthful, budget-feasible, and in large markets achieves an approximation ratio of $\frac{1}{2}$ (with small modifications, this can be converted to a randomized $\frac{1}{2}$ -approximation for arbitrary markets as well Chen et al. (2011)).

3 Our Approach

In this section, we give a high level overview of our approach. Sections 3.1 and 3.2 are preliminary sections and must be read before proceeding further. The notions defined in these sections are explained more intuitively using an example in Section 3.3. Also for rest of the paper we will assume that the sellers' items are divisible, unless we explicitly talk about indivisible items.

3.1 A Notion of An Allocation Rule

To build our new ideas, we first introduce and formalize the notion of an allocation rule.

An allocation rule $f : \mathbb{R}^+ \rightarrow [0, 1]$ is a function which determines how much to buy from a given seller. The domain of allocation rules is the *cost per utility rate*; meaning, given a (u_i, c_i) pair of seller i , the allocation rule f says that we should buy $f(\frac{c_i}{u_i})$ fraction of seller i 's item. We do not enforce using the same allocation rule for all sellers.

We say an allocation rule $f : \mathbb{R}^+ \rightarrow [0, 1]$ is a **Standard Allocation Rule** if f is a decreasing function such that $f(0) = 1$ and $f(e - 1) = 0$.⁵

For any standard allocation rule $f : \mathbb{R}_+ \rightarrow [0, 1]$, we can define an associated family of allocation rules

$$\mathcal{F}(f) = \{f_r : \mathbb{R}_+ \rightarrow [0, 1]\}_{r>0}$$

where f_r denotes an allocation rule which is same as f except that it is stretched along the horizontal axis with ratio r , i.e. $f_r(x) = f(x/r)$ for all $x \geq 0$.

As we will see later, any single standard allocation rule f and its corresponding family of allocation rules $\mathcal{F}(f)$ will uniquely specify our mechanism. At a high level, our mechanism will work as follows: we will pick the largest positive r such that $f_r \in \mathcal{F}(f)$ is budget-feasible; meaning the sum of the payments with allocation rule f_r does not exceed B . However, note that we have not yet defined a payment rule given an allocation rule f_r - we define it next.

⁴ It is worth pointing out that for submodular utilities, they need to use an additional trick: constructing a (sorted) list of sellers in a greedy manner before running the auction.

⁵The choice of $e - 1$ is just for simplifying the future calculations; it can be replaced with any other constant.




Figure 1: With allocation rule f_r , the payment to seller i is defined to be u_i times the shaded area under the curve.

3.2 Payment Rule

Recall that given a function $f_r \in \mathcal{F}(f)$ and (u_i, c_i) pair for a seller i , the value of $f_r(\frac{c_i}{u_i})$ only tells us what fraction of seller i 's item to buy. But how much should we pay seller i in order to give incentives to seller i to report its cost truthfully to the mechanism? We compute these payments based on the well known Myerson's characterization of the truthful mechanisms Myerson (1981).

Let the payment rule for seller i is denoted by $P_{i,r} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ for an allocation rule f_r . Here $P_{i,r}$ maps the reported cost of seller i into its payment.

To define $P_{i,r}$, we do the following thought process: Let's divide seller i 's item into u_i different pieces. Note that now the seller's cost for each piece is $\frac{c_i}{u_i}$. Thus function f_r can now be seen as mapping the cost of a single piece into the fraction of that piece that we will buy. Let $Q_r(x) : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ denote the function that maps the cost for a single piece into a payment for that piece. Now, Myerson's characterization Myerson (1981) says that the payment for each piece is given by the following formula:

$$Q_r(x) = x \cdot f_r(x) + \int_x^\infty f_r(y) \, dy.$$

Intuitively, $Q_r(x)$ represents the area under the curve as seen in Figure 1. Going forward, we will call the function Q_r a *unit-payment rule*. Note that $P_{i,r}(x)$ and Q_r are related by the following formula:

$$P_{i,r}(x) = u_i \cdot Q_r(x/u_i),$$

Thus, to summarize, for an allocation rule f_r , we buy $f_r(c_i/u_i)$ units of her item, and pay her $P_{i,r}(c_i)$ amount of money.

Remark: We make a remark that the above payment rule is truthful only if the allocation rule f_r that is offered to seller i doesn't depend on the private information (cost c_i in this case) of the seller

i. If the allocation rule f_r does depend on the private information of the seller, then the mechanism may or may not be truthful. In the next section we give a mechanism in which the allocation rule f_r for a seller i depends on its reported costs. Later in section 4, we give a mechanism where the allocation rule f_r for a seller i doesn't depend on its reported cost, thus our payment rule will ensure that the resulting mechanism is truthful.

3.3 Example

Suppose the buyer has budget $B = 13/3$, $S = \{s_1, s_2\}$ and sellers s_1, s_2 each own an item with cost 2 and 4, respectively. Also, suppose both of the items have utility 1.

Let the mechanism use the family of curves $\mathcal{F}(f)$ for $f(x) = 1 - x$ where the domain of f is $[0, 1]$. The mechanism should find the largest r for which f_r is budget feasible. To this end, we set r to be a very large number and decrease r until f_r is budget feasible. For instance, suppose we start from $r = 10$ (see Figure 2). Observe that the payment of the mechanism in this case would be 4.8 and 4.2 to s_1 and s_2 , respectively. Since the sum of payments exceed $13/3$, then $r = 10$ is not budget feasible. Consequently, we reduce r further until the mechanism becomes budget feasible at $r = 4$: At this point, the payment of the mechanism to s_1 and s_2 would respectively be $8/3$ and $5/3$, which sum up to be exactly B . (see Figure 3)




Figure 2: The mechanism is not budget feasible at $r = 10$.




Figure 3: The mechanism is budget feasible at $r = 6$.

3.4 First Attempt: A Parameterized Class of Envy-Free Mechanisms

In this section we describe a mechanism (denoted by Mechanism Envy-Free(f)) that is not always truthful, but it will form the basis of our truthful mechanism. Moreover, some structural results about this mechanism will be useful while analyzing our truthful mechanism, thus we will be talking about this mechanism throughout the paper. This mechanism is parameterized by the choice of a standard allocation rule f . The mechanism described in this section offers a single allocation rule $f_r \in \mathcal{F}(f)$ to all the sellers, thus it is envy-free (although it may not be truthful).

Definition 2 We say that an allocation rule f_r is a budget-feasible allocation rule if $\sum_{i \in S} P_{i,r}(c_i) = B$, i.e. the payments defined with respect to f_r sum up to B .

Now given any standard allocation rule f , the mechanism starts with a very large scaling ratio $r = \infty$ so that we are guaranteed to have $\sum_{i \in S} P_{i,r}(c_i) > B$.

Then, the mechanism decreases r until the rule f_r becomes a budget-feasible rule (say at $r = r^*$). The mechanism stops at this point and uses f_{r^*} and $\{P_{i,r^*}\}_{i \in S}$ to determine the allocations and payments. The ratio r^* is also called the *stopping rate* of the mechanism. We define this process formally in Mechanism Envy-Free(f).

Mechanism Envy-Free: Parameterized by a standard allocation rule f

input : Budget B , (u_i, c_i) pair for each seller i

output: A scaling ratio r^*

$r \leftarrow \infty$;

while f_r is not a budget-feasible rule **do**

 | Decrease r slightly;

end

$r^* \leftarrow r$;

Output the scaling ratio r^* ;

One can easily see that the above mechanism is budget-feasible, individually rational, and envy-free; however, it may not be truthful. Also, the efficiency of the above mechanism depends on the choice of function f . Thus, an important question is: What is the optimal choice of function f ? Let's first understand the performance of the above mechanism for a simple choice of function f .

Definition 3 A standard allocation rule $f : \mathbb{R}^+ \rightarrow [0, 1]$ is called a uniform standard allocation rule if $f(x) = 1$ for $x < e - 1$, and $f(x) = 0$ otherwise. Figure 4 depicts this curve.

One can show that the above envy-free mechanism when run using a uniform standard allocation rule, mimics the simple factor $\frac{1}{2}$ mechanism presented earlier. Thus, it turns out to be truthful as well for this choice of standard allocation rule. However for more general allocation rules, the above envy-free mechanism might not be truthful. Thus before we answer the harder question about the optimal choice of function f , we next describe the truthful version of the above envy-free mechanism.




Figure 4: The Uniform Standard Allocation Rule

4 A Parameterized Class of Truthful Mechanisms

We use a simple trick to convert Mechanism Envy-Free(f) to a truthful mechanism. The idea is to define, for each seller i , an allocation rule which does not depend on c_i . In particular, we define the allocation rule for seller i to be f_{r_i} , where r_i will be chosen independently of c_i . For finding r_i , we run Mechanism Envy-Free(f) on the instance which is obtained by setting c_i to be 0 while keeping cost of the other sellers intact; r_i would be the stopping rate of the mechanism Envy-Free(f). The formal definition of the truthful mechanism appears in Mechanism Truthful(f).

Mechanism Truthful(f): Parameterized by a standard allocation rule f

input : Budget B , (u_i, c_i) pair for each seller i

foreach $i \in S$ **do**

$temp \leftarrow c_i$;
 $c_i \leftarrow 0$;
 $r_i \leftarrow \text{Mechanism Envy-Free}(f)$;
 $c_i \leftarrow temp$;

end

foreach $i \in S$ **do**

Allocate $f_{r_i}(c_i)$ from seller i ;
 Pay $P_{i,r_i}(c_i)$ to seller i ;

end

In Lemma 2, we prove that Mechanism Truthful(f) is individually rational, truthful, and budget-feasible for any given standard allocation rule f . First, we state the following useful lemma.

Lemma 1 For any seller $i \in S$ we have $r^* \geq r_i$.

Proof. The proof is based on the fact that $P_{i,r}(x)$ is an increasing function of r (for a fixed x) and is a decreasing function of x (for a fixed r). The proof is by contradiction, suppose $r^* < r_i$. Let

$c'_j = c_j$ for all $j \in S \setminus \{i\}$ and let $c'_i = 0$. Observe that

$$\begin{aligned} B &= \sum_{j \in S} P_{j,r^*}(c_j) \leq \sum_{j \in S} P_{j,r^*}(c'_j) \\ &< \sum_{j \in S} P_{j,r_i}(c'_j) \end{aligned}$$

where the first inequality is due to the fact that $P_{j,r^*}(x)$ is a decreasing function of x and the second inequality is due to the fact that $r^* < r_i$. However, note that the above inequalities imply that $B < \sum_{j \in S} P_{j,r_i}(c'_j)$, which contradicts with the budget feasibility of Mechanism Envy-Free(f): see that $\sum_{j \in S} P_{j,r_i}(c'_j)$ represents the payment of Envy-Free(f) when the costs are c'_1, \dots, c'_n , and so it can not be larger than B . \square

Lemma 2 *Mechanism Truthful(f) is individually rational, truthful, and budget-feasible.*

Proof. Note that the allocation and payment rules for seller i , i.e. f_{r_i}, P_{i,r_i} , do not depend on the cost reported by her. This fact, along with the fact that f_{r_i} is a monotone rule (decreasing function) implies individual rationality and truthfulness. The proof is almost identical to the proof of Myerson's Lemma and we do not repeat it here.

The proof for budget feasibility needs a bit more work. Let p_i, p'_i denote the payments to seller i respectively in Mechanism Truthful(f) and Mechanism Envy-Free(f), i.e. $p_i = P_{i,r_i}(c_i)$ and $p'_i = P_{i,r^*}(c_i)$. The lemma is proved if we show that $p_i \leq p'_i$, since we have $\sum_{i \in S} p'_i = B$.

To see $p_i \leq p'_i$, note that $P_{i,r}(x)$ is an increasing function of r (for a fixed x). So, since we have $r^* \geq r_i$ due to Lemma 1, it must be the case that $P_{i,r_i}(c_i) \leq P_{i,r^*}(c_i)$. \square

5 A $(1 - 1/e)$ -Approximate Optimal Truthful Mechanism

So far, we have introduced a parameterized class of individually rational, truthful, and budget-feasible mechanisms for the problem: Passing any standard allocation rule f to Mechanism Truthful(f) fixes the mechanism which we denote by \mathcal{M}_f . Our goal in this section is to find the *most efficient* mechanism in this class. Formally, given a standard allocation rule f , we denote the approximation ratio of \mathcal{M}_f by \mathcal{R}_f and define it as:

$$\mathcal{R}_f = \inf_I \frac{\mathcal{U}_f(I)}{\mathcal{U}^*(I)},$$

where the infimum is taken over all instances I of the problem ⁶. Here $\mathcal{U}_f(I)$ denotes the utility obtained by \mathcal{M}_f in instance I , and $\mathcal{U}^*(I)$ denotes the optimum utility in instance I .

⁶If we are focused on large markets, we take only instances I for which the largeness ratio is smaller than some threshold, and take the limiting approximation factor as the threshold goes to 0.

The most efficient allocation rule f , is the one which maximizes \mathcal{R}_f . Our goal, in this section and Section 6, is to find the most efficient allocation rule and its corresponding approximation ratio. Formally, we prove the following theorem.

Theorem 1 *The most efficient standard allocation rule for Mechanism Truthful(f) is $f(x) = \ln(e - x)$, for which we get $\mathcal{R}_f = 1 - 1/e$, i.e. it has an approximation ratio $1 - 1/e$.*

We prove this theorem in two parts: In the first part we show that $\mathcal{R}_f \geq 1 - 1/e$ for $f(x) = \ln(e - x)$; this is proved in the current section. In the second part, we show that $\mathcal{R}_g \leq 1 - 1/e$ for any standard allocation rule g . This fact can be seen as a consequence of our hardness result in Section 6, which states that no truthful mechanism can achieve approximation ratio better than $1 - 1/e$. We also provide a more direct (alternative) proof in Section D that shows our choice of $f(x) = \ln(e - x)$ is optimal among all possible choices of the standard allocation rules.

5.1 Finding an optimal f for the (non-truthful) Mechanism Envy-Free(f)

In this section, we prove that Mechanism Envy-Free(f) has approximation ratio $1 - 1/e$ for $f(x) = \ln(e - x)$. Note that the Mechanism Envy-Free(f) is not truthful, however its analysis will be helpful when analyzing our truthful mechanism in Section 5.2 and in Section A. Here, we analyze Mechanism Envy-Free(f) assuming that the true costs are known; later, in Section 5.2, we use this result to prove that Mechanism Truthful(f) has approximation ratio $1 - 1/e$ for the same choice of f .

5.1.1 Preliminaries

We use g_r to denote the inverse of an allocation rule f_r , i.e. $g_r(x) = f_r^{-1}(x)$. Given an allocation rule f_r , we also write an alternative definition of its corresponding unit-payment rule Q_r . This definition, rather than being in terms of $\frac{c_i}{u_i}$, would be in terms of $f_r(\frac{c_i}{u_i})$. This alternative definition is denoted by G_r , and is defined such that $Q_r(\frac{c_i}{u_i}) = G_r(f_r(\frac{c_i}{u_i}))$. For instance, if a seller owns an item with utility 1, then we pay her $G_r(x)$ when a fraction x of her item is allocated. To be more precise, for $y = f_r(\frac{c_i}{u_i})$ we define

$$G_r(y) = \int_0^y g_r(x) dx = Q_r\left(\frac{c_i}{u_i}\right).$$

We also denote g_1 and G_1 respectively by g and G .

Proposition 1 *Given the standard allocation rule $f(x) = \ln(e - x)$, it is straight-forward to verify that $g(x) = e - e^x$ and $G(x) = ex - e^x + 1$. Also, $f_r(x) = \ln\left(\frac{er-x}{r}\right)$.*

From now on in this section, we assume that $f(x) = \ln(e - x)$. Next, we prove a useful inequality in the following lemma which will be used in the analysis of Envy-Free(f).

Lemma 3 For any x, α such that $0 \leq x, \alpha \leq 1$ we have

$$G(x) - \alpha \cdot g(x) \leq e \cdot (x - \alpha \cdot (1 - 1/e)).$$

Proof.

$$\begin{aligned} \alpha(e^x - 1) &\leq e^x - 1 \\ \Rightarrow \alpha e^x - e^x + 1 &\leq \alpha \\ \Rightarrow \alpha e^x - e^x + 1 + e(x - \alpha) &\leq \alpha + e(x - \alpha) \\ \text{(by the definition of } g, G) \Rightarrow G(x) - \alpha \cdot g(x) &\leq e \cdot (x - \alpha \cdot (1 - 1/e)). \end{aligned}$$

□

5.1.2 Approximation Ratio of Mechanism Envy-Free(f)

In the following lemma, we prove the efficiency of Mechanism Envy-Free(f) when all sellers report true costs.

Lemma 4 If sellers report true costs, then Mechanism Envy-Free(f) has approximation ratio $1 - 1/e$.

Proof. Observe that w.l.o.g. we can assume $r^* = 1$: If $r^* \neq 1$, then we can construct a new instance which is *similar* to the original instance and has stopping rate 1. More precisely, there exists some $\beta > 0$ such that if we multiply the budget and the reported costs by β , the stopping rate becomes equal to 1. Note that this operation will not change the optimal solution or the solution of Envy-Free(f) and can be performed w.l.o.g.

Now, suppose that a fraction x_i of item i is allocated by Envy-Free(f). Since $r^* = 1$, we can use Lemma 3 to write the following set of inequalities:

$$G(x_i) - \alpha_i \cdot g(x_i) \leq e \cdot (x_i - \alpha_i \cdot (1 - 1/e)) \quad \forall i \in S,$$

where α_i is the fraction that is allocated from seller i in the optimal solution (recall that we are comparing Envy-Free(f) with the optimum fractional solution). The above inequalities can be multiplied by u_i on both side and be written as:

$$u_i \cdot (G(x_i) - \alpha_i \cdot g(x_i)) \leq u_i \cdot e \cdot (x_i - \alpha_i \cdot (1 - 1/e)) \quad \forall i \in S.$$

By adding up these inequalities, we get:

$$\sum_{i \in S} u_i \cdot (G(x_i) - \alpha_i \cdot g(x_i)) \leq e \cdot \sum_{i \in S} u_i \cdot (x_i - \alpha_i \cdot (1 - 1/e)). \quad (1)$$

Now, we show that if

$$0 \leq \sum_{i \in S} u_i \cdot (G(x_i) - \alpha_i \cdot g(x_i)), \quad (2)$$

then the lemma is proved using (1) and (2). First we show why (1) and (2) prove the lemma, and then in the end, we prove (2) itself.

Observe that (1) and (2) imply that

$$0 \leq \sum_{i \in S} u_i \cdot (x_i - \alpha_i \cdot (1 - 1/e)). \quad (3)$$

Now, let U denote the utility gained by Envy-Free(f) and $U^* = \sum_{i \in S} u_i \alpha_i$ denote the utility of the optimum (fractional) solution; see that (3) implies

$$(1 - 1/e) \cdot U^* = \sum_{i \in S} \alpha_i u_i \cdot (1 - 1/e) \leq \sum_{i \in S} x_i u_i = U,$$

This would prove the lemma.

So, it only remains to show that (2) holds: First observe that $\sum_{i \in S} u_i \cdot G(x_i) = B$, since the sum represents the payment of Envy-Free(f). Also, see that $\sum_{i \in S} \alpha_i u_i \cdot g(x_i) \leq B$, since this sum is a lower bound on the cost of the optimal solution, which is at most B . \square

5.2 Special case: Analyzing our truthful mechanism for unit utilities

In this section, we prove that Mechanism Truthful(f) has approximation ratio $1 - 1/e$ in large markets when it uses the standard allocation rule $f(x) = \ln(e - x)$ for the special case when all the utilities are equal to 1. In other words, we will show that approximation ratio approaches $1 - 1/e$ as θ , the market's largeness ratio, approaches 0 for the case of unit utilities. The proof for the case of general utilities is intricate and appears in Section A.

Note that the assumption of unit utilities imply $c_i/u_i = c_i$ for any seller i . Next, we state two lemmas before proving the approximation ratio. For simplicity in the analysis, w.l.o.g., assume that $c_1 \leq c_2 \leq \dots \leq c_n$.

Lemma 5 $r_1 \geq r_2 \geq \dots \geq r_n$.

Lemma 6 Let $u^*(b) : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ denote the maximum utility that the buyer can achieve with budget b (when the items are divisible). Then, $u^*(b)$ is a concave function.

Proofs for both of these lemmas are straight-forward and are deferred to the appendix, Section B.

Lemma 7 *Mechanism Truthful(f) has approximation ratio $1 - 1/e$ when all the items have utility equal to 1.*

Proof. Recall that $U^* = u^*(B)$ and let U denote the utility achieved by Mechanism Truthful(f). We need to show that $(1 - 1/e) \cdot U^* \leq U$. Instead of showing that $U = \sum_{i \in S} f_{r_i}(c_i)$ is large enough compared to U^* , we show that $\sum_{i \in S} f_{r_n}(c_i)$ is large enough compared to U^* ; the lemma then would be proved since we have $f_{r_n}(c_i) \leq f_{r_i}(c_i)$ for all $i \in S$. To see why $f_{r_n}(c_i) \leq f_{r_i}(c_i)$, it is enough to note that $r_n \leq r_i$ due to Lemma 5 which implies $f_{r_n}(c_i) \leq f_{r_i}(c_i)$.

We consider two cases for the proof: In Case 1 we assume $c_{\max} \leq \bar{c}$, and in Case 2 we assume otherwise, where the number \bar{c} is the cost at which $f_{r_n}(\bar{c}) = 1 - 1/e$, more precisely, this happens at $\bar{c} = r_n(e - e^{1-1/e})$.

Case 1 In this case, observe that we have $f_{r_n}(c_i) \geq 1 - 1/e$ for all $i \in S$, which implies $f_{r_i}(c_i) \geq 1 - 1/e$. This just means $U \geq (1 - 1/e)n \geq (1 - 1/e)U^*$.

Case 2 Let $U_n = \sum_{i \in S} f_{r_n}(c_i)$, we will show that

$$U_n \geq (1 - 1/e) \cdot (1 - o(1)) \cdot U^*. \quad (4)$$

To prove this, consider an auxiliary instance in which, instead of budget B , we have a reduced budget $B' = \sum_{i \in S} Q_{r_n}(c_i)$. Note that if we run Mechanism Envy-Free(f) on the auxiliary instance, then its stopping rate is r_n , and so, the utility gained by the mechanism is exactly U_n . Let U_{aux}^* denote the optimal utility in the auxiliary instance. Then, by applying Lemma 4 on the auxiliary instance, we have $U_n \geq (1 - 1/e) \cdot U_{\text{aux}}^*$. So, if we show that

$$U_{\text{aux}}^* \geq (1 - o(1)) \cdot U^* \quad (5)$$

then (4) holds and the proof is complete.

We use Lemma 6 to prove (5): First, we show that $B' \geq (1 - o(1)) \cdot B$; then, applying Lemma 6 would imply that $u^*(B') \geq (1 - o(1)) \cdot u^*(B)$, which is identical to (5) by definition. So all we need to complete the proof is showing that $B' \geq (1 - o(1)) \cdot B$.

To this end, we prove that $B' \geq (1 - \alpha \cdot \frac{c_{\max}}{B}) \cdot B$, where α is a constant with value $(e - e^{1-1/e})^{-1} \approx 6/5$. This would prove the Lemma due to the large market assumption. First, observe that

$$\begin{aligned} B &= Q_{r_n}(0) + \sum_{i \in S \setminus \{n\}} Q_{r_n}(c_i) \leq Q_{r_n}(0) + B' \\ \Rightarrow B' &\geq B - Q_{r_n}(0) \geq B - r_n. \end{aligned} \quad (6)$$

Now, recall that in Case 2, we have $c_{\max} \geq \bar{c}$, which implies

$$\begin{aligned} B &\geq \bar{c} \cdot \frac{B}{c_{\max}} = r_n(e - e^{1-1/e}) \cdot \frac{B}{c_{\max}} \\ &\Rightarrow c_{\max} \cdot (e - e^{1-1/e})^{-1} \geq r_n. \end{aligned} \tag{7}$$

Combining (6) and (7) implies $B' \geq (1 - \alpha \cdot \frac{c_{\max}}{B}) \cdot B$ with the promised value for α . \square

6 Impossibility Result: On why $1 - 1/e$ is the best approximation possible

In this section we show that no truthful (and possibly) randomized mechanism achieves approximation ratio better than $1 - 1/e$. We prove a stronger claim by allowing the mechanism to satisfy the budget feasibility in expectation, i.e. we prove that no truthful mechanism that is budget-feasible in expectation can achieve a ratio better than $1 - 1/e$. From now on, in this section, we assume that all the mechanisms that we refer to are truthful, and are also budget-feasible in expectation. First, we prove the claim assuming that the items are indivisible; then we will see that the same proof easily extends to divisible items as well.

Proof Outline. We construct a bayesian instance of the problem and prove that no budget-feasible truthful mechanism for this instance can achieve approximation ratio better than $1 - 1/e$; this also implies that no mechanism for the prior-free setting can achieve ratio better than $1 - 1/e$ ⁷. The proof is done in two steps. First, we show that for any truthful mechanism for this instance, there exists a simple posted price mechanism that achieves at least the same revenue. The posted price mechanism simply offers the same price p to every seller and pays p to any seller who accepts the offer and 0 to others. In the second step of the proof, we show that for any choice of p , such mechanisms cannot achieve a ratio better than $1 - 1/e$. The proof that we present, w.l.o.g. analyzes the market in expectation: budget feasibility is satisfied in expectation; also, the utility of the mechanisms are computed in expectation.

We now give the full proof by first giving our hardness instance.

The Impossibility Instance. We construct a bayesian instance of the problem in which all the sellers have unit utility and their costs are drawn i.i.d. from a distribution with cumulative

⁷This is so because an α -approximate mechanism in the prior-free setting is also α -approximate in Bayesian setting

distribution function F , defined as follows:

$$F(x) = \begin{cases} 1/e & \text{if } x = 0, \\ \frac{1}{e(1-x)} & \text{if } 0 < x \leq 1 - 1/e. \end{cases}$$

In other words, $F(x)$ denotes the probability that the cost of a seller is at most x . Let \mathcal{D} be the distribution defined by F and let \bar{c} denote the expected cost of a seller sampled from \mathcal{D} , i.e. $\bar{c} = \mathbb{E}_{x \sim \mathcal{D}}[x]$. We define the budget to be $B = \bar{c} \cdot N$ where $N \geq 1$ is an arbitrary integer denoting the number of sellers.

Definition 4 *A posted price mechanism is a mechanism that offers a price p_i to any seller $i \in S$, and pays her p_i if she accepts the offer and pays her 0 otherwise.*

Definition 5 *A uniform posted price mechanism is a posted price mechanism that offers the same price to all sellers.*

Definition 6 *A cut-off allocation rule is an allocation rule which allocates the whole unit of an item if its cost is less than a certain cut-off and allocate 0 units otherwise. Let $\text{cutoff}(p)$ denote a cut-off allocation rule with the cut-off price p .*

It is clear that posted price mechanisms use cut-off allocation rules to allocate items from sellers.

Lemma 8 *If the sellers costs are drawn i.i.d. from the distribution \mathcal{D} , then for any mechanism in this bayesian setting there exists a posted price mechanism with the same approximation ratio.*

Proof. Due to Myerson's Lemma Myerson (1981), any truthful mechanism in Bayesian setting can be seen as a set of allocation and payment rules corresponding to each seller, where the allocation rule is a decreasing function (of the cost) and the payment rule is defined with respect to the allocation rule as we saw in Figure 1. Given such an allocation rule for an arbitrary seller $i \in S$, namely A_i , one can think of a simpler way to implement A_i (in expectation) by finding a distribution π_i over cut-off allocation rules.

More precisely, the distribution π_i will be a distribution over cut-off prices, which can also be interpreted in the following way: We first sample a price x according to π_i , and then present that price to seller i , i.e. run $\text{cutoff}(x)$.

Let the CDF of π_i be denoted by $F_i(\cdot)$ and let its PDF be denoted by $f_i(\cdot)$. We prove the existence of π_i in the following claim.

Claim 1 *Define the distribution π_i by its CDF $F_i(\cdot)$ such that $F_i(x) = 1 - A_i(x)$ for any cost $x \geq 0$. Then π_i would implement A_i in expectation.*

Proof. All we need to show is that a seller with price x would be allocated with probability $A_i(x)$ in π_i . To see this, note that the probability that the seller is allocated is exactly equal to the probability of observing a cut-off price at least x when a cut-off price is sampled from π_i . This probability is equal to $1 - (1 - A_i(x))$ by the definition of π_i ; this proves the claim. \square

Now, we claim that the cut-off allocation rule with the cut-off price

$$p_i = F^{-1} \left(\int_0^\infty f_i(p) \cdot F(p) dp \right) \quad (8)$$

achieves the same utility and spends the same budget (in expectation) as the allocation rule A_i paired with its corresponding Myerson payment rule.

Claim 2 For any seller $i \in S$, $\text{cutoff}(p_i)$ achieves the same utility and spends the same budget (in expectation) as the allocation rule A_i paired with its corresponding Myerson payment rule.

Proof. The main idea of the proof is that the set of points $P = \{(F(x), xF(x)) : 0 \leq x < 1\}$ forms a straight line in the two-dimensional plane; See Figure 5 for a proof by picture. Note that $F(x)$




Figure 5: This parametric plot is representing the set of points P for $x = 0$ to $x = 1$. The horizontal and vertical axis respectively represent $F(x)$ and $xF(x)$

denotes the expected allocation when a price x is offered to a seller and $xF(x)$ is the corresponding expected payment.

Now, see that the expected utility achieved by the allocation rule A_i is $\int_0^\infty f_i(p)F(p)dp$, which is exactly equal to the expected utility achieved by $\text{cutoff}(p)$ due to (8). To prove the claim, it remains to verify that the Myerson payment rules corresponding to $\text{cutoff}(p_i)$ and A_i spend the same budget (in expectation). To this end, just observe that P is a straight line; consequently, since the allocation rules A_i and $\text{cutoff}(p_i)$ allocate equal units of items (in expectation), they also spend the same amount of the budget (in expectation). \square

Due to Claim 2, the posted price mechanism that offers price p_i to seller i is budget-feasible (in expectation) and also achieves an expected utility equal to the utility of the originally given mechanism. This proves the claim. \square

Lemma 9 *If the sellers costs are drawn i.i.d. from the distribution \mathcal{D} , then for any (budget-feasible) posted price mechanism there exists a (budget-feasible) uniform posted price mechanism with the same approximation ratio.*

Proof. Suppose that $\{p_i\}_{i \in S}$ denotes the offered prices in a posted price mechanism and let

$$\bar{p} = F^{-1} \left(\frac{1}{|S|} \cdot \sum_{i \in S} F(p_i) \right).$$

First, observe that the uniform posted price mechanism with price \bar{p} achieves a utility equal to the utility of the original posted price mechanism; this is simply because $F(p)$, the expected utility gained by offering a price of p , is the average of $F(p_i)$ for $i = 1, \dots, N$. It remains to verify that the uniform posted price mechanism is budget-feasible. To this end, just observe that the set of points P (depicted in Figure 5) is a straight line; consequently, since the posted price mechanism and the uniform posted price mechanism allocate equal units of items (in expectation), they also spend the same amount of the budget (in expectation). \square

Theorem 2 *For the case of indivisible items, no truthful budget-feasible mechanism can achieve approximation ratio better than $1 - 1/e$.*

Proof. We use Lemma 9 and show that no uniform posted price mechanism can achieve ratio better than $1 - 1/e$. Equivalently, we show that the uniform posted price mechanism which spends all the budget in expectation has approximation ratio no better than $1 - 1/e$.

The uniform posted price mechanism that spends all the budget in expectation offers a price p^* such that $p^* F(p^*) \cdot N = B$. Given the definitions of $F(\cdot)$ and B , we can solve this equation to get $p^* = \frac{e-2}{e-1}$. Now, we are ready to compute the approximation ratio. First, note that the (expected) utility of the uniform posted price mechanism is $N \cdot F(p^*)$. If it was true that $\sum_{i \in S} c_i \leq B$, then we would have $U^* = N$ (the optimum solution could buy all items), and so we could write the approximation ratio as

$$\frac{N \cdot F(p^*)}{N} = F(p^*) = 1 - 1/e,$$

which would prove the claim. However, although the condition is satisfied in expectation, i.e. $\mathbb{E} [\sum_{i \in S} c_i] = B$, the sum is not always bounded by B , which means $U^* = N$ does not always hold. We find a way to fix this issue using Hoeffding bounds (see Section F to see formal statements of

Hoeffding bounds). We show that although $\sum_{i \in S} c_i$ is not always bounded by B , it is concentrated around its mean, B , with high probability. We will see that this is enough to prove the theorem.

As a consequence of Hoeffding bounds (stated in Section F), for any $\epsilon > 0$ we have:

$$\Pr \left[\sum_{i \in S} c_i \geq (1 + \epsilon) \cdot B \right] \leq e^{-\Omega(|S|)} \quad (9)$$

Recall that we defined $N = |S|$ and that in our hardness instance $N \rightarrow \infty$. Using (9), we will provide an upper bound on the approximation ratio which, for any constant $\epsilon > 0$, approaches to $(1 - 1/e)(1 + \epsilon)$ as N approaches infinity. This proves that the approximation ratio can not be a constant larger than $1 - 1/e$.

To this end, first note that if $\sum_{i \in S} c_i \leq B(1 + \epsilon)$, then we have $U^* \geq \frac{N}{1 + \epsilon} - 1$; this holds due to Lemma 6. We can use this fact along with (9) to write the following upper bound on the (expected) approximation ratio:

$$(1 - e^{-\Omega(N)}) \cdot \frac{N \cdot F(p^*)}{N/(1 + \epsilon) - 1} + e^{-\Omega(N)} \cdot 1.$$

The above ratio clearly approaches $F(p^*)(1 + \epsilon)$ as $N \rightarrow \infty$. Noting that $F(p^*) = 1 - 1/e$ finishes the proof. \square

Now we use Theorem 2 to prove its counterpart for divisible items.

Corollary 1 *For the case of divisible items, no truthful budget-feasible mechanism can achieve approximation ratio better than $1 - 1/e$.*

Proof. Proof by contradiction. Suppose there exists a mechanism with approximation ratio $\alpha > 1 - 1/e$ for some constant α . Then, we show that we can convert this mechanism to an α -approximation mechanism for indivisible items which is truthful and budget-feasible in expectation. This would contradict Theorem 2.

To do this conversion, we repeat the exact same argument that we used to prove Theorem 2. As the result, we can convert the given α -approximation mechanism to a uniform posted price mechanism with approximation ratio α . Note that all posted price mechanisms allocate items without dividing them. Consequently, we have an α -approximation mechanism for indivisible items. Contradiction. \square

7 Conclusion

Our main contribution is designing optimal budget-feasible mechanisms for the knapsack model in large markets. First, we assume that the items are divisible, and study a natural class of deterministic mechanisms: each mechanism in this class is characterized by a decreasing allocation

function. All the mechanisms in this class are individually rational, truthful and budget-feasible, but they have different approximation ratios based on the choice of the allocation function. We find a mechanism in this class which has an approximation ratio $1 - 1/e$, and prove that no truthful mechanism can achieve a better approximation ratio.

We also provide a mechanism with approximation ratio $1 - 1/e$ for the case of indivisible items: the idea is to first run the mechanism for divisible items, and then round the obtained fractional solution (allocation). We design a rounding process that takes the fractional allocation as its input and outputs an integral allocation with its associated payments. Due to the properties of our rounding process, the resulting mechanism is individual rational, truthful-in-expectation, and budget-feasible; also, it has approximation ratio $1 - 1/e$ in large markets.

Finally, we study the problem for submodular utility functions with indivisible items. For this case, we first design a deterministic mechanism which has approximation ratio $\frac{1}{2}$ in large markets; this mechanism can have an exponential running time in general. Inspired by this mechanism, we also design a polynomial-time deterministic mechanism with approximation ratio $\frac{1}{3}$. We do not provide any results for when the items are divisible in the submodular model: One has to model the utility function over divisible items; the multilinear extension Vondrak (2008) or Lovász extension of submodular functions is a potential choice for this purpose. We leave open this case for future study. Bayesian

Acknowledgments

We acknowledge the comments and ideas of an anonymous reviewer that helped us generalize our impossibility result to the case of bayesian setting.

References

- Amazon’s mechanical turk platform. <https://www.mturk.com/>.
- Clickworker, virtual workforce. <http://www.clickworker.com>.
- Crowdfunder. <http://www.crowdfunder.com/>.
- Hoeffding bounds. http://en.wikipedia.org/wiki/Hoeffding's_inequality.
- BADANIDIYURU, A., KLEINBERG, R., AND SINGER, Y. 2012. Learning on a budget: Posted price mechanisms for online procurement. In *EC*.
- BEI, X., CHEN, N., GRAVIN, N., AND LU, P. 2012. Budget feasible mechanism design: from prior-free to bayesian. In *STOC*.

- CHEN, N., GRAVIN, N., AND LU, P. 2011. On the approximability of budget feasible mechanisms. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '11. SIAM, 685–699.
- DEVANUR, N. R. AND HAYES, T. P. 2009. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *ACM Conference on Electronic Commerce*.
- DOBZINSKI, S., PAPADIMITRIOU, C. H., AND SINGER, Y. 2011. Mechanisms for complement-free procurement. In *ACM Conference on Electronic Commerce*. 273–282.
- FELDMAN, J., HENZINGER, M., KORULA, N., MIRROKNI, V. S., AND STEIN, C. 2010. Online stochastic packing applied to display ad allocation. In *ESA (1)*. 182–194.
- FELDMAN, J., KORULA, N., MIRROKNI, V. S., MUTHUKRISHNAN, S., AND PÁL, M. 2009. Online ad assignment with free disposal. In *WINE*. 374–385.
- GOEL, G. AND MEHTA, A. 2008. Online budgeted matching in random input models with applications to adwords. In *SODA*.
- HOREL, T., IOANNIDIS, S., AND MUTHUKRISHNAN, S. 2013. Budget feasible mechanisms for experimental design. *CoRR abs/1302.5724*.
- MEHTA, A., SABERI, A., VAZIRANI, U. V., AND VAZIRANI, V. V. 2007. Adwords and generalized online matching. *J. ACM* 54, 5.
- MYERSON, R. B. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 58–73.
- SINGER, Y. 2010. Budget feasible mechanisms. In *FOCS*. 765–774.
- SINGER, Y. 2011. How to win friends and influence people, truthfully: Influence maximization mechanisms for social networks. In *WSDM*.
- SINGER, Y. AND MITTAL, M. 2013. Pricing mechanisms for crowdsourcing markets. In *WWW*. 1157–1166.
- SINGLA, A. AND KRAUSE, A. 2013a. Incentives for privacy tradeoff in community sensing. In *AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*.
- SINGLA, A. AND KRAUSE, A. 2013b. Truthful incentives in crowdsourcing tasks using regret minimization mechanisms. In *WWW*. WWW '13. 1167–1178.
- SVIRIDENKO, M. 2004. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters* 32, 1, 41 – 43.

VONDRAK, J. 2008. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*. STOC '08. ACM, 67–74.

A Analyzing Our Optimal Truthful Mechanism for General Additive Utilities (Continuation of Section 5)

In this section we will prove that the approximation ratio of mechanism $\text{Truthful}(f)$ approaches $1 - 1/e$ as θ , the market's largeness ratio, approaches 0. We emphasize that here we dismiss the extra assumption that was made in Section 5.2: There, we assumed all items have utility 1, here we give a proof for the general case when item i provides utility $u_i > 0$.

Lemma 10 *For each $k \in \{1, \dots, n\}$, $r_k \geq (1 - \theta)r^*$.*

Proof. We just need to prove that $f_{(1-\theta)r^*}$ is not a budget-tight rule (i.e. does not consume all of the budget) when we set the cost of item k to 0. First of all, note that

$$Q_{(1-\theta)r^*}(x) = (1 - \theta)Q_{r^*}\left(\frac{x}{1 - \theta}\right) \leq (1 - \theta)Q_{r^*}(x)$$

Here we used the fact that Q_{r^*} is a decreasing function. This implies that $\sum_{i=1}^n u_i Q_{(1-\theta)r^*}\left(\frac{c_i}{u_i}\right) \leq (1 - \theta)B$. This expression is the budget consumed by the rule $f_{(1-\theta)r^*}$ without setting the cost of item k to 0. When we set c_k to 0, the amount of budget consumed can be bounded in the following manner

$$u_k Q_{(1-\theta)r^*}(0) + \sum_{i \neq k} u_i Q_{(1-\theta)r^*}\left(\frac{u_i}{c_i}\right) \leq (1 - \theta)B + u_k \left(Q_{(1-\theta)r^*}(0) - Q_{(1-\theta)r^*}\left(\frac{c_k}{u_k}\right) \right) \quad (10)$$

Note that $Q_{(1-\theta)r^*}(\cdot)$ is defined to be the area of the shaded region as seen in figure 1. Therefore one can crudely upper bound the difference $Q_{(1-\theta)r^*}(0) - Q_{(1-\theta)r^*}(x)$ by $x \times f_{(1-\theta)r^*}(0)$ for any $x \geq 0$. Now letting $x = \frac{c_k}{u_k}$, and substituting in inequality 10 we get

$$\begin{aligned} u_k Q_{(1-\theta)r^*}(0) + \sum_{i \neq k} u_i Q_{(1-\theta)r^*}\left(\frac{u_i}{c_i}\right) &\leq (1 - \theta)B + u_k \left(\frac{c_k}{u_k} - 0\right) \\ &= (1 - \theta)B + c_k \leq B \end{aligned}$$

This completes the proof. □

Lemma 11 *Mechanism $\text{Truthful}(f)$ has an approximation ratio approaching $1 - \frac{1}{e}$ as θ approaches 0.*

Proof. W.l.o.g. assume that $r^* = 1$ (since we can scale the budget and costs by an appropriate scaling factor). Now let us pick a constant threshold $0 < s < e - 1$ and partition the indices $\{1, \dots, n\}$ into two sets \mathcal{I} and \mathcal{J} : let \mathcal{J} be the set of indices i where $\frac{c_i}{u_i} > s$ and let \mathcal{I} be the complement.

Let r^+ be the minimum r_i where $i \in \mathcal{J}$. If \mathcal{J} happens to be empty, let $r^+ = r^* = 1$. Let B' be the budget consumed by the allocation rule f_{r^+} , i.e. let $B' = \sum_{i=1}^n u_i Q_{r^+}(\frac{c_i}{u_i})$. We will prove that B' is close to B . If $r^+ = r^*$, this is obviously true because $B' = B$. So assume that $r^+ = r_k$ for some $k \in \mathcal{J}$.

Because of the way r_k is chosen, we have

$$B = u_k Q_{r_k}(0) + \sum_{i \neq k} u_i Q_{r_k}(\frac{c_i}{u_i}) \leq u_k + \sum_{i \neq k} u_i Q_{r_k}(\frac{c_i}{u_i}) \quad (11)$$

Here we used the fact that $Q_{r_k}(0) \leq Q_{r^*}(0) = 1$ (since we assumed $r^* = 1$). Note that $B' \geq \sum_{i \neq k} u_i Q_{r^+}(\frac{c_i}{u_i})$. Combining this with the inequality 11 we get

$$B' \geq B - u_k = B - c_k \frac{u_k}{c_k} \geq B - \frac{c_k}{s} \geq (1 - \frac{\theta}{s})B$$

Using lemma 6, one can see that $u^*(B') \geq (1 - \frac{\theta}{s})u^*(B)$. But we also know from lemma 4 that the utility achieved by f_{r^+} is at least $(1 - \frac{1}{e})u^*(B')$. Therefore we have

$$\sum_{i=1}^n u_i f_{r^+}(\frac{c_i}{u_i}) \geq (1 - \frac{\theta}{s})(1 - \frac{1}{e})u^*(B) \quad (12)$$

For an item $i \in \mathcal{I}$, we have $r_i \geq (1 - \theta)r^* = 1 - \theta$ (we used lemma 10). Therefore

$$\frac{f_{r_i}(\frac{c_i}{u_i})}{f(\frac{c_i}{u_i})} = \frac{f(\frac{1}{r_i} \frac{c_i}{u_i})}{f(\frac{c_i}{u_i})} \geq \frac{f(\frac{1}{1-\theta} \frac{c_i}{u_i})}{f(\frac{c_i}{u_i})}$$

One can easily verify that $\ln f$ is a concave function. Therefore $\frac{f(\frac{1}{r_i}x)}{f(x)}$ for $x \leq s$ is minimized at $x = s$. This means that

$$\frac{f_{r_i}(\frac{c_i}{u_i})}{f(\frac{c_i}{u_i})} \geq \frac{f(\frac{s}{r_i})}{f(s)} \geq \frac{f(\frac{s}{1-\theta})}{f(s)}$$

If we let $\alpha = \frac{f(\frac{s}{1-\theta})}{f(s)}$, then for every $i \in \mathcal{I}$ we have

$$f_{r_i}(\frac{c_i}{u_i}) \geq \alpha f(\frac{c_i}{u_i}) \geq \alpha f_{r^+}(\frac{c_i}{u_i})$$

Similarly, for every item $i \in \mathcal{J}$, $r_i \geq r^+$ and therefore $f_{r_i}(\frac{c_i}{u_i}) \geq f_{r^+}(\frac{c_i}{u_i}) \geq \alpha f_{r^+}(\frac{c_i}{u_i})$.

We just proved that for every $i \in \{1, \dots, n\}$, $f_{r_i}(\frac{c_i}{u_i}) \geq \alpha f_{r^+}(\frac{c_i}{u_i})$. Combining this with inequality 12 we get

$$\sum_{i=1}^n u_i f_{r_i}(\frac{c_i}{u_i}) \geq \alpha (1 - \frac{\theta}{s})(1 - \frac{1}{e}) u^*(B)$$

So the approximation ratio for Mechanism Truthful(f) is at least

$$\alpha (1 - \frac{\theta}{s})(1 - \frac{1}{e}) = \frac{\ln(e - \frac{s}{1-\theta})}{\ln(e-s)} (1 - \frac{\theta}{s})(1 - \frac{1}{e})$$

For any fixed s , strictly smaller than $e - 1$, one can observe that the ratio above approaches $1 - \frac{1}{e}$ as $\theta \rightarrow 0$. We will not attempt to optimize the value of s for the sake of brevity. \square

B Missing Proofs from Section 5.2

Proof of Lemma 5:

For any i, j such that $i \leq j$, we prove that $r_i \geq r_j$, this would prove the lemma. The proof is by contradiction, suppose $r_i < r_j$. First, note that $Q_{r_i}(0) + \sum_{k \in S \setminus \{i\}} Q_{r_i}(c_k)$ represents the sum of payments in Envy-Free(f) when the cost of i is set to 0. Since Envy-Free(f) spends all of the budget, then we have:

$$Q_{r_i}(0) + \sum_{k \in S \setminus \{i\}} Q_{r_i}(c_k) = B. \quad (13)$$

Also, note that $Q_{r_j}(0) + \sum_{k \in S \setminus \{j\}} Q_{r_j}(c_k)$ represents the sum of payments in Envy-Free(f) when the cost of j is set to 0; a similar argument shows that

$$Q_{r_j}(0) + \sum_{k \in S \setminus \{j\}} Q_{r_j}(c_k) = B. \quad (14)$$

Taking the difference between (13) and (14) implies that

$$\left(Q_{r_i}(0) - Q_{r_j}(0) \right) + \left(Q_{r_i}(c_j) - Q_{r_j}(c_i) \right) = 0 \quad (15)$$

Now, observe that since $r_i < r_j$ and $c_i \leq c_j$, then we have that $Q_{r_i}(0) < Q_{r_j}(0)$ and $Q_{r_i}(c_j) < Q_{r_j}(c_i)$. This contradicts with (15). \square

Proof of Lemma 6: We need to prove that for $B_1, B_2 \geq 0$ and $0 \leq \lambda \leq 1$

$$\lambda u^*(B_1) + (1 - \lambda) u^*(B_2) \leq u^*(\lambda B_1 + (1 - \lambda) B_2)$$

Let x_i be the amount of item i we allocate to achieve $u^*(B_1)$ and let y_i be the amount of item

i we allocate to achieve $u^*(B_2)$.

Now let $z_i = \lambda x_i + (1 - \lambda)y_i$. Note that since $0 \leq x_i, y_i \leq 1$, we also have $0 \leq z_i \leq 1$. If we allocate z_i from item i , the utility we get will be

$$\sum_{i=1}^n u_i z_i = \lambda \left(\sum_{i=1}^n u_i x_i \right) + (1 - \lambda) \left(\sum_{i=1}^n u_i y_i \right) = \lambda u^*(B_1) + (1 - \lambda) u^*(B_2)$$

The cost paid by these allocations is simply

$$\sum_{i=1}^n c_i z_i = \lambda \left(\sum_{i=1}^n c_i x_i \right) + (1 - \lambda) \left(\sum_{i=1}^n c_i y_i \right) \leq \lambda B_1 + (1 - \lambda) B_2$$

Therefore z_i 's are an allocation that spend a budget of at most $\lambda B_1 + (1 - \lambda) B_2$ and yet achieve a utility of $\lambda u^*(B_1) + (1 - \lambda) u^*(B_2)$. This proves that

$$u^*(\lambda B_1 + (1 - \lambda) B_2) \geq \lambda u^*(B_1) + (1 - \lambda) u^*(B_2)$$

□

C Mechanisms for Indivisible Items

In this section, using our mechanism for divisible items, we design a mechanism for indivisible items with approximation ratio $1 - 1/e$. The idea is to first run the mechanism for divisible items, and then round the obtained fractional solution (allocation). We design a rounding process that takes the fractional allocation as its input and outputs an integral allocation with its associated payments. Due to the properties of our rounding process, the resulting mechanism is individual rational, truthful, and budget-feasible; also, it has approximation ratio $1 - 1/e$ in large markets.

First, we explain a set of properties that we need the rounding procedure to satisfy. If the rounding procedure satisfies these properties, then its individual rationality, truthfulness, and budget feasibility would be guaranteed. Also, these properties guarantee that the approximation ratio would remain $1 - 1/e$. First we explain these properties in Section C.1, then, we state our rounding procedure and prove that it satisfies these desired properties in Section C.2.

C.1 Properties of the Rounding Procedure

Let $\tilde{x}_1, \dots, \tilde{x}_n$ represent a fractional allocation where \tilde{x}_i denotes the allocated fraction from seller $i \in S$; also, let $\tilde{p}_1, \dots, \tilde{p}_n$ be the associated payments for this allocation. We round this solution to an integral solution, represented by the allocation x_1, \dots, x_n and payments p_1, \dots, p_n , such that:

1. Item i is bought with probability \tilde{x}_i .

2. If item i is bought, then $p_i = \tilde{p}_i/\tilde{x}_i$, and $p_i = 0$ otherwise.
3. $\sum_{i \in S} p_i \leq B + c_{\max}$.

Properties 1 and 2 imply individual rationality and truthfulness: Verifying individual rationality is straight-forward due to the individual rationality of the fractional solution. For truthfulness, just see that $\mathbb{E}[x_i] = \tilde{x}_i$ and $\mathbb{E}[p_i] = \tilde{p}_i$, which implies $\mathbb{E}[p_i - c_i x_i] = \tilde{p}_i - c_i \tilde{x}_i$. This just means that, in the mechanism for indivisible items, i cannot benefit (in expectation) by misreporting, since she is already receiving the maximum possible expected utility that she can ever achieve.

Properties 1 and 2 also imply budget feasibility in expectation, however, Property 3 provides a much stronger guarantee: the budget will not be violated by an additive factor more than c_{\max} .⁸

C.2 Description of the Rounding Procedure

In this section, we focus in designing a rounding procedure which satisfies Properties 1, 2 and 3. To this end, we first need to define a polytope \mathcal{P} that represents all the (fractional) allocations which, in a certain sense, are budget-feasible:

$$\mathcal{P} = \left\{ y \in [0, 1]^n : \sum_{i \in S} y_i \cdot \frac{\tilde{p}_i}{\tilde{x}_i} \leq B \right\}$$

First, we prove that extreme points of \mathcal{P} are “almost” integral.

Definition 7 *A point $y \in [0, 1]^n$ is called semi-integral if there is at most one entry of y which is non-integral, i.e. there is at most one index i such that $0 < y_i < 1$.*

Lemma 12 *All the extreme points of \mathcal{P} are semi-integral.*

Proof. The proof is straight-forward, we give a high-level description and omit the formal details. The idea is to see \mathcal{P} as the intersection of a hypercube and a hyperplane; the hypercube is $[0, 1]^n$ and the hyperplane is $\sum_{i \in S} y_i \cdot \frac{\tilde{p}_i}{\tilde{x}_i} \leq B$. So, any extreme point of \mathcal{P} either is an extreme point of the hypercube (which is integral), or is on the intersection of the hyperplane and an edge of the hypercube. In the latter case, it can be seen that such a point has at most one fractional entry, since any two adjacent vertices on the hypercube are different in at most one entry. \square

Outline of the Rounding Procedure The procedure accepts the fractional allocation constructed by the mechanism, i.e. $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_n)$, and then writes it as a convex combination of extreme points of \mathcal{P} . Then, it samples an extreme point from the convex combination, where each point is selected with probability proportional to its coefficients in the convex combination. Finally,

⁸We can always reduce the budget slightly to get strict budget feasibility, e.g. we can reduce the budget to $(1 - \epsilon)B$ for an arbitrary small $\epsilon > 0$. This will not affect the approximation ratio (asymptotically) in a large market.

it rounds the sampled extreme point (which is a semi-integral point) to an integral point. We use the following fact about semi-integral points for implementing the last step:

Fact 1 *A semi-integral point $y \in [0, 1]^n$ can be written as the convex combination of two integral points which differ in at most one entry, i.e. $y = \alpha y' + (1 - \alpha)y''$ where $y', y'' \in \{0, 1\}^n$ are integral points which differ in at most one entry.*

Now we are ready to formally state our main rounding procedure.

Procedure EfficientRounding

input : Allocation vector \tilde{x} and Payment vector \tilde{p}

- 1 Find extreme points $z^1, \dots, z^K \in \mathcal{P}$ and positive numbers $\lambda_1, \dots, \lambda_K$ summing up to one such that $\tilde{x} = \sum_{i=1}^K \lambda_i \cdot z^i$;
 - 2 Sample a single point from $\{z^1, \dots, z^K\}$ where z^i is selected with probability λ_i ; Let z denote the sampled point;
 - 3 Write z as the convex combination of two integral points x^1, x^2 such that x_1, x_2 differ in at most one entry, i.e. suppose $z = \alpha \cdot x^1 + (1 - \alpha)x^2$;
 - 4 With probability α , let $x = x^1$, otherwise, let $x = x^2$;
 - 5 Announce x as the final allocation and pay $x_i \cdot \tilde{p}_i / \tilde{x}_i$ to seller i .
-

Lemma 13 *Procedure EfficientRounding satisfies Properties 1, 2 and 3.*

Proof. It is straight-forward to verify that Properties 1 and 2 hold; for any seller $i \in S$ we have:

$$\begin{aligned} \mathbb{E}[x_i] &= \mathbb{E}[\alpha \cdot x_i^1 + (1 - \alpha) \cdot x_i^2] \\ &= \mathbb{E}[z_i] = \mathbb{E} \sum_{j=1}^K \lambda_j \cdot z_i^j = \tilde{x}_i \end{aligned}$$

which implies Property 1 since x_i is a binary random variable. Property 2 trivially holds by the construction of Procedure EfficientRounding.

It remains to prove Property 3. To this end, define $p(y) = \sum_{i \in S} y_i \cdot \tilde{p}_i / \tilde{x}_i$ for any $y \in [0, 1]^n$. We prove the claim by showing that $p(x) \leq B + c_{\max}$. Equivalently, we can show that $p(x^1) \leq B + c_{\max}$ and $p(x^2) \leq B + c_{\max}$. We prove this only for x^1 , the proof for x^2 is identical. The claim is trivial if $x^1 = x^2$, since in this case we have $x^1 \in \mathcal{P}$, which means $p(x^1) \leq B$. So suppose $x^1 \neq x^2$. Recall that we have

$$z = \alpha \cdot x^1 + (1 - \alpha)x^2,$$

where x^1, x^2 are two adjacent vertices on the hypercube. Also, recall that any two adjacent vertices on the hypercube are different in exactly one entry, so suppose x^1, x^2 are different in entry j , i.e.

$x_j^1 \neq x_j^2$. Now, we prove the lemma by showing that

$$p(x_1) \leq p(z) + \tilde{p}_j \cdot \frac{1 - \tilde{x}_j}{\tilde{x}_j} \leq B + c_{\max}. \quad (16)$$

First verify that the first inequality in (16) holds since x_1 and z are only different in their j -th entry: if $x_j^1 = 0$, then $p(x_1) \leq p(z)$; if $x_j^1 = 1$, then it is straight-forward to verify that $p(x_1) = p(z) + \tilde{p}_j(1 - \tilde{x}_j)/\tilde{x}_j$. Having that the first inequality holds, (16) is proved if we show that

$$p(z) \leq B, \quad (17)$$

$$\tilde{p}_j(1 - \tilde{x}_j)/\tilde{x}_j \leq c_{\max}. \quad (18)$$

To verify (17), just note that $z \in \mathcal{P}$. To verify (18), recall that $\tilde{x}_j = f_r(c_j/u_j)$ and $\tilde{p}_j = u_j \cdot Q_r(c_j/u_j)$ for some $r > 0$; we will prove the following bounds on r :

$$\tilde{p}_j/\tilde{x}_j \leq ru_j, \quad (19)$$

$$r \leq \frac{c_j/u_j}{1 - \tilde{x}_j}. \quad (20)$$

Observe that combining (19) and (20) implies (18). So, we are done if we prove (19) and (20) hold. To prove (19), it is enough to note that $Q_r(c_j/u_j) \leq r\tilde{x}_j$, i.e. the area under the curve f_r that represents the payment per unit of utility to seller j fits in a rectangle with width r and height \tilde{x}_j ; this implies $\tilde{p}_j/\tilde{x}_j \leq ru_j$. It is also straight-forward to verify (20) holds due to the concavity of f_r . \square

D The Optimal Standard Allocation Rule

Mechanism Truthful(f) is defined uniquely by the standard allocation rule f . Here we provide an alternative proof for showing that our choice of $f(x) = \ln(e - x)$ is optimal, in the sense that Mechanism Truthful(f) attains the best approximation ratio under this choice. Although this fact is also a consequence of what we already proved in Section 6, we provide a more direct proof here with a prior-free hardness instance.

For simpler analysis, we first ignore the trick that we used for making our mechanism truthful: we work with the non-truthful, but cleaner mechanism Envy-Free(f) and we prove that no approximation ratio better than $1 - \frac{1}{e}$ is attainable using any standard allocation rule f . In the end, we see that the trick that makes mechanisms truthful only worsens the approximation ratio. Therefore the optimality result also applies to the family of truthful mechanisms defined by Mechanism Truthful(f).

We provide hardness examples with unit utilities, i.e. with $u_i = 1$ for all $i \in S$. We also

drop the indices from payment functions P . This is done because all the payment functions are identical: Note that $\text{Envy-Free}(f)$ uses the same allocation rule for all the sellers, which means the unit-payment functions are identical. Also, since all the utilities are 1, then the unit-payment functions are identical to the payment functions, i.e. $P(c) = Q(c)$. Unless specified otherwise we also assume that the scaling ratio r^* is 1. i.e. we have $P(c) = Q(c) = Q_1(c)$ and $f(c) = f_1(c)$.

Definition 8 *Assume that we are given a standard allocation rule f . Define the set $S_f \subseteq \mathbb{R}^2$ as follows*

$$S_f = \{(f(c), P(c) - c) \mid c \in \mathbb{R}\}$$

where $P(c)$ is the payment rule associated with the allocation rule f (i.e. $P(c) = Q_1(c)$). Also define the set T_f as the downward closure of S_f , i.e.

$$T_f = \{(x, y) \in \mathbb{R}^2 \mid \exists y' \geq y : (x, y') \in S_f\}$$

Note that in our definition we simply work with the standard function f and not its scaled variants f_r . The following lemma gives us a way to find out if the approximation ratio attained by the mechanism is worse than a number β .

Lemma 14 *Suppose that the point $(\beta, 0)$ lies in the convex hull of T_f . Then the mechanism defined by f has approximation ratio at most β .*

Proof. Assume that the point $(\beta, 0)$ can be written as the convex combination of points in T_f in the following way

$$(\beta, 0) = \sum_{i=1}^n \alpha_i p_i$$

where $\alpha_i \geq 0$, $\sum_{i=1}^n \alpha_i = 1$, and $p_1, \dots, p_n \in T_f$. Because of the way T_f is defined, for each p_i one can find a value c_i such that $p_i = (f(c_i), y_i)$ and $y_i \leq P(c_i) - c_i$. Then this means that

$$\sum_{i=1}^n \alpha_i f(c_i) = \beta \tag{21}$$

$$\sum_{i=1}^n \alpha_i (P(c_i) - c_i) \geq 0 \tag{22}$$

Now let M be a very large number and consider M sellers. For each i , let α_i fraction of the sellers price their item at c_i . The issue of $\alpha_i M$ not being integral can be easily dealt with, but would unnecessarily complicate the proof, hence we assume $\alpha_i M$ is integral. Assume that each seller's item is worth one unit to the buyer. Now let the buyer's budget B be

$$B = M \sum_{i=1}^n \alpha_i P(c_i)$$

With this definition, we are simply saying that f is budget-tight with respect to our budget, or in other words, the mechanism consumes all of the budget at $r = 1$.

Because of (22), our budget B is at least as large as the sum of the costs of all items $M \sum_{i=1}^n \alpha_i c_i$. Therefore we could buy all of the items and obtain a utility of M , if we were not constrained by the mechanism. But because of (21), our mechanism obtains a utility of βM . Therefore the approximation ratio in this example is $\frac{\beta M}{M} = \beta$. \square

It is worth mentioning that the reverse of lemma 14 is also true under some mild conditions, but in this section we are only concerned with the direction proved.

Building on top of our lemma, we can now prove that $1 - \frac{1}{e}$ is the best approximation factor among the mechanisms defined by a standard allocation rule.

Theorem 3 *Given any standard allocation rule f , the mechanism $Envy-Free(f)$ has approximation ratio at most $1 - \frac{1}{e}$.*

Proof. To prove this, we simply need to show that the point $(1 - \frac{1}{e} + \epsilon, 0)$ is inside the convex hull of T_f for every $\epsilon > 0$.

Suppose the contrary is true. So for some $\epsilon > 0$, the point $(\beta, 0) = (1 - \frac{1}{e} + \epsilon, 0)$ is not inside the convex hull of T_f . Since the point is not inside the convex hull, there is a line that separates them. We can further assume that this line passes through the point and is therefore given by the equation $y = s(x - \beta)$ for some slope s . Since T_f is downward closed with respect to the y coordinate, T_f must fall completely below this line. Since $S_f \subseteq T_f$, this means that S_f must also fall completely below this line which means that for all c , we have

$$P(c) - c \leq s(f(c) - \beta) \tag{23}$$

We can safely assume that $s > 0$. If this was not true, then for $c = 0$, we would necessarily have $f(c) < \beta$. But this already shows that the mechanism never buys more than a fraction of β of any item, and therefore the approximation factor cannot be larger than β .

Now consider the function $f_0(x) = \ln(\max(e - tx, 1))$ for some given $t > 0$. For this allocation function let P_0 be the Myerson's payment rule, i.e. $P_0(x) = xf_0(x) + \int_x^\infty f_0(x)dx$ (the shaded area in figure 1). One can easily verify that for $0 \leq c \leq \frac{e-1}{t}$

$$\int_c^\infty f_0(x)dx = c - \frac{e-1}{t} - (c - \frac{e}{t}) \ln(e - tc)$$

and therefore

$$P_0(c) - c = \frac{e}{t} \ln(e - tc) - \frac{e-1}{t} = \frac{e}{t} (f_0(c) - (1 - \frac{1}{e})) \tag{24}$$

If we choose t so that $s = \frac{e}{t}$, then equation 24 becomes very similar to the inequality 23. If we

subtract equation 24 from the inequality 23, then we get

$$P(c) - P_0(c) \leq s(f(c) - f_0(c) - \epsilon) \quad (25)$$

This inequality holds for $0 \leq c \leq \frac{\epsilon-1}{t} = s(1 - \frac{1}{e})$. For $c = 0$, we get that $P(0) - P_0(0) \leq s(f(0) - 1 - \epsilon) < 0$. Therefore $P(0) < P_0(0)$.

Let $c^* = \sup\{c \mid f(c) < f_0(c)\}$. This supremum is strictly greater than 0, because otherwise we would have $f(x) \geq f_0(x)$ for all $x > 0$, which would mean that $P(0) \geq P_0(0)$, which is a contradiction. This supremum is also finite because for $c \geq \frac{\epsilon-1}{t}$ we have $f_0(c) = 0$ and $f(c) \geq 0$.

Because c^* is the supremum, we can find a sequence of points c_1, c_2, \dots , such that $f(c_i) < f_0(c_i)$ for all i , and $c_i \rightarrow c^*$. For each such c_i , we have $c_i \leq c^* \leq \frac{\epsilon-1}{t}$, and therefore the inequality 25 holds. This means that

$$c_i(f(c_i) - f_0(c_i)) + \int_{c_i}^{\infty} (f(x) - f_0(x))dx \leq s(f(c_i) - f_0(c_i) - \epsilon)$$

By rearranging the terms we get

$$\int_{c_i}^{\infty} (f(x) - f_0(x))dx \leq (s - c_i)(f(c_i) - f_0(c_i)) - s\epsilon \leq -s\epsilon$$

Here we used the fact that $c_i \leq c^* \leq s(1 - \frac{1}{e}) < s$ and $f(c_i) - f_0(c_i) < 0$.

But now as we take the limit as $i \rightarrow \infty$, we get

$$\int_{c^*}^{\infty} (f(x) - f_0(x))dx \leq -s\epsilon$$

which is a contradiction because for $x > c^*$, we have $f(x) - f_0(x) \geq 0$, and so the integral cannot be negative. \square

Theorem 4 *For any standard allocation rule f , Mechanism Truthful(f) has approximation ratio at most $1 - \frac{1}{e}$.*

Proof. Mechanism Truthful(f) always achieves a worse or equal utility compared to Envy-Free(f). Therefore its approximation ratio is at most $1 - \frac{1}{e}$ by Theorem 3. \square

E Submodular Utility Functions

In this section, we present truthful mechanisms for the knapsack problem when the utility function for the buyer is a *monotone* submodular function rather than an additive function. More precisely, a monotone submodular function $F : 2^S \rightarrow \mathbb{R}^+$ defines the utility that the buyer derives from buying

a subset of S . The buyers problem then becomes selecting a subset $S^* \subset S$ such that S^* is budget feasible, i.e. $c(S^*) \leq B$, and S^* has the highest utility, $F(S^*)$, among all the budget feasible subsets. (recall that $c(S^*)$ denotes $\sum_{i \in S^*} c_i$)

This problem has first been studied in Singer (2010) for arbitrary markets and a 0.0089-approximation is presented for it. Later, Chen et al. (2011) improved this result by giving an exponential-time deterministic mechanism with approximation ratio 0.119 and a polynomial-time randomized mechanism with approximation ratio 0.126.

Our Results In this section, we focus on studying this problem in *large markets* where each individual can not significantly affect the market (see Section E.1 for a formal definition). Under this assumption, we design a deterministic mechanism which has approximation ratio $\frac{1}{2}$, but has an exponential running time. Later, we will see that the exponential running time is solely due to the computational difficulty of solving the knapsack problem for submodular functions. In fact, our mechanism is also a polynomial-time $(\gamma^2/2)$ -approximation when it has access to a γ -approximation oracle for solving the knapsack problem (see Section E.3). To the extent of our knowledge, the best existing approximation oracle has $\gamma = 1 - 1/e$ due to Sviridenko (2004); this provides us a mechanism with approximation ratio $\gamma^2/2 \approx 0.2$.

We take a step further and improve this result by presenting a deterministic polynomial-time $\frac{1}{3}$ -approximation mechanism in Section E.4. This mechanism, although using a greedy optimization oracle with $\gamma = 1 - 1/e$, has approximation ratio equal to $\frac{1}{3}$ (rather than $\gamma^2/2 \approx 0.2$).

Oneway-Truthfulness All of the mechanisms that we have are truthful, however, we first present a simpler version of them which are not fully truthful but satisfy truthfulness in a weaker form, which we call *oneway-truthfulness*. Briefly, by this property, players only have incentive to report costs lower than their true cost. This notion is formally defined in Section E.2.2.

In Section E.5, we convert our oneway-truthful mechanisms to (fully) truthful mechanisms only by changing the payment rule. It is worth pointing out that analyzing the performance ratio of oneway-truthful mechanisms is not much different than analyzing truthful mechanisms: Since the cost of optimum solution may only decrease if players report lower costs, then any α -approximate solution for the reported instance is also an α -approximate solution for the original instance.

E.1 Preliminaries

In this section, first we state a few basic definitions. Then, we formally define the large market assumption and oneway-truthfulness. Finally, we state a few definitions regarding submodular functions which are used in our mechanisms.

E.2 Basic Definitions

Similar to before, we say a subset of sellers $T \subseteq S$ is budget feasible if $c(T) \leq B$. Utility of the subset T is defined by $F(T)$. The *optimum subset*, S^* , is the budget feasible subset with the highest utility. We call $F(S^*)$ the *optimum utility* and also denote it by F^* .

E.2.1 The Large Market Assumption

Our large market assumption here is almost identical to the alternative large market assumption that was discussed in Section 1.1.1. Intuitively, it says that no individual affects the (optimum solution of the) market significantly. This assumption is formally defined below.

Let $u_{\max} = \max_{s \in S} F(\{s\})$ and U^* be the total utility of the optimum solution (i.e. the maximum achievable utility when the costs are known). This large market assumption states that

Definition 9 *We say that a market is large if $u_{\max} \ll U^*$.*

In other words, we define the largeness ratio of the market to be $\theta = \frac{u_{\max}}{U^*}$ and analyze our mechanisms for when $\theta \rightarrow 0$.

E.2.2 Oneway-Truthfulness

Think of a reverse auction with a set of sellers S where each seller $i \in S$ has a private cost c_i . In a truthful mechanism, no seller wants to report a fake cost regardless of what others do. In a oneway-truthful mechanism, no seller wants to report a cost higher than its true cost regardless of what others do.

For clarification, we first define the notion of cost vector briefly: when we say a cost vector d , we mean a vector which has an entry d_i corresponding to any seller i , where d_i represents the cost associated with seller i . Now we formally define the notion of oneway-truthfulness as follows:

Definition 10 *A mechanism \mathcal{M} is oneway-truthful if, for any seller $i \in S$ and any cost vector d for which $d_i > c_i$, we have:*

$$u_i(c_i, d_{-i}) \geq u_i(d)$$

where d_{-i} denotes any cost vector corresponding to the rest of players except i and $u_i(\cdot)$ denotes the utility of player i .

E.2.3 Submodular Functions

Given the submodular function $F : 2^S \rightarrow \mathbb{R}^+$, we define an ordering of the elements of S with respect to F , which we call the *greedy sequence* and denote it by $\chi(F) = \langle x_1, \dots, x_n \rangle$. For simplicity

in the definition, we first define an auxiliary notion as follows: let $\chi_i = \cup_{j=1}^i \{x_j\}$ for all i , and let $\chi_0 = \emptyset$. The sequence is constructed such that

$$x_i = \arg \max_{s \in S \setminus \chi_{i-1}} F(\chi_{i-1} \cup \{s\})$$

for all positive $i \leq n$.

It is easy to verify that $\chi(F)$ can be constructed in polynomial time by finding the values of x_1, \dots, x_n one by one in the order that they appear in $\chi(F)$. After constructing of the greedy sequence, define $\partial_i = F(\chi_i) - F(\chi_{i-1})$ for all $i \leq n$.

E.3 The Exp-Time Mechanism

In this Section, we present an extremely simple mechanism which we call the *Oracle Mechanism*. Given a submodular function $F : 2^S \rightarrow \mathbb{R}^+$, the mechanism first finds the optimal budget feasible subset, i.e. the subset $S^* \subseteq S$ such that S^* is budget feasible and has the highest utility among all the budget feasible subsets. Let $F^* = F(S^*)$ and $r^* = \frac{B}{F^*}$. We also call F^*, r^* respectively the *optimum utility* and the *optimum cost per utility rate*. We also call r^* the *optimum rate* when there is no risk of confusion.

Winner Selection The Mechanism constructs the sequence $\chi(F)$ and chooses the largest integer k such that $F(\chi_k) \leq F^*/2$. Then, it reports χ_k as the set of winners.

The Payment Rule For simplicity, assume that the winners are indexed from $1, \dots, k$. The Payment to winner i is equal to $2r_i \cdot \partial_i$, where r_i is the optimum cost per utility rate for the instance in which seller i is removed from the set of sellers, S . In other words, think of an auxiliary instance in which we are given budget B and the cost of every seller is the same as the original instance except that $c_i = \infty$. Then, r_i is the optimum cost per utility rate in this instance.

Below we prove that the Oracle Mechanism is individually rational, oneway-truthful, and it has approximation ratio $\frac{1}{2}$ in large markets. Also, it is *almost budget feasible*, i.e. we can show that the total sum of its payments is at most $B + o(B)$. So, all we need for having a strictly budget feasible mechanism is starting with a slightly decreased budget. We will see that this does not affect the approximation ratio of the mechanism asymptotically due to the large market assumption. The formal proof for this is deferred to Section E.6.

Simplifying Assumption Through out the analysis, w.l.o.g. we assume that the sellers appear in the greedy sequence in an increasing order, i.e. $x_i = i$ for all $i \in S$.

Lemma 15 *The Oracle Mechanism is individually rational.*

Mechanism 1: Oracle Mechanism

input : Submodular utility function F , Budget B

Sellers report their costs;

Compute the optimum utility F^* and the optimum cost per utility rate r^* ;

Construct the sequence $\chi(F)$;

Find the largest integer k such that $F(\chi_k) \leq F^*/2$;

Announce χ_k as the set of winners;

foreach $i \in \chi_k$ **do**

$S \leftarrow S \setminus \{i\}$;

 Let r_i be the optimum cost per utility rate in the current instance;

 Pay $2r_i \cdot \partial_i$ to seller i ;

$S \leftarrow S \cup \{i\}$;

end

Proof. We show that $c_i \leq 2r^* \cdot \partial_i$ and $r^* \leq r_i$. These two imply $c_i \leq 2r_i \cdot \partial_i$ which is individual rationality. First we prove $r^* \leq r_i$ as follows. Let F_i^* denote the optimum utility when seller i is removed. Clearly, we have $F_i^* \leq F^*$. This just means $r^* \leq r_i$ since we have $r_i F_i^* = r^* F^* = B$.

It remains to show that $c_i \leq 2r^* \cdot \partial_i$. Note that we only need to show this for the last winner, i.e. it suffices to prove that $c_k \leq 2r^* \cdot \partial_k$ since we have $\frac{c_i}{\partial_i} \leq \frac{c_j}{\partial_j}$ iff $i \leq j$. The proof is by contradiction, suppose $c_k > 2r^* \cdot \partial_k$.

For more intuition, we first explain our argument for contradiction in words and then we state it more formally. From the definition of the greedy sequence $\chi(F)$, it can be seen that conditioned on buying the subset χ_k , the cost for buying each extra unit of utility is at least $\frac{c_k}{\partial_k}$, which is more than $2r^*$. So, even if we get the subset χ_k for free, the cost for buying an additional $F(S^*) - F(\chi_k)$ units of utility (which is needed for the optimum solution) would be

$$\begin{aligned} (F(S^*) - F(\chi_k)) \cdot \frac{c_k}{\partial_k} &\geq \frac{F^*}{2} \cdot \frac{c_k}{\partial_k} \\ &> \frac{F^*}{2} \cdot 2r^* = B. \end{aligned}$$

This means cost of the optimum solution is more than B .

To formalize this contradiction, just note that by monotonicity of F , we have

$$F(S^* \cup \chi_k) - F(\chi_k) \geq \frac{F^*}{2}. \quad (26)$$

Now observe that by the definition of the greedy sequence $\chi(F)$, the cost for buying each extra unit of utility conditioned on having χ_k is at least $\frac{c_k}{\partial_k}$. This fact, and (26) together imply that

$$c(S^* \cup \chi_k) - c(\chi_k) \geq \frac{c_k}{\partial_k} \cdot \frac{F^*}{2}. \quad (27)$$

On the other hand, recall that $c_k > 2r^* \cdot \partial_k$, so we can write (27) as

$$c(S^* \cup \chi_k) - c(\chi_k) \geq \frac{c_k}{\partial_k} \cdot \frac{F^*}{2} > r^* \cdot F^* = B.$$

This implies $c(S^*) > B$ which is a contradiction with the budget feasibility of S^* . \square

Lemma 16 *The Oracle Mechanism is oneway-truthful.*

Proof. For contradiction, suppose there exists a seller $i \in S$ who has incentive to report a cost \bar{c}_i which is higher than its true cost c_i . Assuming that the mechanism picked χ_k as the set of winners, we have either $i > k$ or $i \leq k$. The proof is done separately in each of these cases.

If $i > k$, then see that seller i can not change the first $i - 1$ elements of $\chi(F)$ by reporting a higher cost. Now, see that if $i > k + 1$, then again χ_k will be chosen as the set of winners, which is a contradiction with the incentive of seller i for misreporting. If $i = k + 1$, then see that χ_k will be chosen as a subset of the winners; in this case, the set of winners can possibly contain other sellers, but not certainly not seller $k + 1$ (due to the monotonicity of F). Consequently, seller i remains a loser even by reporting \bar{c}_i . Contradiction.

It remains to do the proof for when $i \leq k$. Recall that the payment to seller i is $2r_i \cdot \partial_i$. Since r_i is not a function of the cost reported by seller i , then see that the only way that i can increase her utility by misreporting, is increasing ∂_i . But by reporting a cost higher than c_i , seller i can not change the first $i - 1$ elements of $\chi(F)$, which means she can only decrease ∂_i by reporting a higher cost. So, the payment to i does not increase if she reports a higher cost. This concludes the lemma. \square

Lemma 17 *In a θ -large market, the Oracle Mechanism has approximation ratio $\frac{1}{2} - \theta$, i.e. asymptotically equal to $\frac{1}{2}$ in a large market.*

Proof. It is enough to note that $F(\chi_{k+1}) \geq F^*/2$, which implies $F(\chi_k) \geq F^* \cdot (1/2 - \theta)$ due to the large market assumption. \square

Now, we prove that sum of the payments in the Oracle Mechanism is at most $B + o(B)$, or in simple words, it is *almost budget feasible*.

Definition 11 *A mechanism is almost budget feasible if its payments sum up to at most $B + o(B)$.*

As we mentioned before, (in large markets) we can convert any almost budget feasible mechanism to a budget feasible mechanism without any loss in its approximation ratio (asymptotically). This can be done simply by running the mechanism with a slightly reduced budget; the proof is deferred to Section E.6.

Lemma 18 *The Oracle Mechanism is almost budget feasible.*

Proof. Recall that the sum of payments is equal to $\sum_{i=1}^k r_i \cdot \partial_i$. To prove the lemma, it is enough to show that for any seller $i \in S$, we have $r_i \leq r^* \cdot (1 - \theta)^{-1}$; because then we have

$$\sum_{i=1}^k r_i \cdot \partial_i \leq r^* \cdot (1 - \theta)^{-1} \cdot \sum_{i=1}^k \partial_i = r^* \cdot (1 - \theta)^{-1} \cdot F(\chi_k) \leq B(1 - \theta)^{-1}.$$

which proves the lemma.

To prove $r_i \leq r^* \cdot (1 - \theta)^{-1}$, let F_i^* denote the optimum utility when seller i is removed. By the large market assumption, we have $\frac{F_i^*}{F^*} \geq 1 - \theta$. This fact, and the fact that $r_i F_i^* = r^* F^* = B$, imply that $r_i \leq \frac{r^*}{1 - \theta}$. \square

The Oracle Mechanism in Polynomial Time

In the Oracle Mechanism, we solve a submodular optimization problem which cannot be solved in polynomial-time, i.e. finding the optimum cost per utility rate which is equivalent to finding the optimum budget feasible subset. Although this problem cannot be solved in polynomial-time, there are approximation algorithms that can find near-optimal solutions for it.

Definition 12 *Suppose we are given an instance of the problem with a submodular function F , budget B , and (publicly known) costs and utilities. A polynomial-time algorithm for solving this problem is called a γ -approximation oracle if, for any instance, it finds a solution with utility at least $\gamma \cdot F^*$.*

Given a γ -approximation oracle, we can run the Oracle mechanism in polynomial time by finding estimates for F^* and r_i 's using the γ -approximation oracle, i.e. instead of computing F^* and r_i 's directly, we compute them using the γ -approximation oracle. The following theorem clarifies the resulting mechanism further and states the properties it satisfies.

Theorem 5 *Suppose the Oracle Mechanism has access to a polynomial-time γ -approximation oracle for the knapsack optimization problem. Then, it is a polynomial-time $(\gamma/2)$ -approximation mechanism and its payments sum up to at most $(B + o(B))/\gamma$.*

Proof. We use the γ -approximation oracle for finding an estimate (lower bound) for F^* . Instead of computing F^* directly, let F^* be the solution which is returned by the γ -approximation oracle. Also, for computing r_i , remove seller i and then compute the optimal utility, namely F_i^* , using the γ -approximation oracle. Everything else remains identical to Mechanism 1.

The proofs for individual rationality and truthfulness follow from the proofs for Mechanism 1. It remains to prove that the payments sum up to at most $B \cdot (1 + o(1))/\gamma$. For this, we follow the proof of Lemma 18. The key point in the proof of Lemma 18 was that $r_i \leq r^* \cdot (1 - \theta)^{-1}$. Here, we prove a weaker inequality $r_i \leq r^* \cdot (1 - \theta)^{-1}/\gamma$. Given this inequality, the rest of the proof remains

similar to the proof of Lemma 18. We do not repeat the full proof here and just show that the weaker inequality holds.

Recall that in this proof, F^* and F_i^* denote the utilities computed by the γ -approximation oracle. Then, by the large market assumption, we have $F_i^* \geq F^* \cdot (1 - \theta)\gamma$. This fact, and the fact that $r_i F_i^* = r^* F^* = B$, imply that $r_i \leq r^* \cdot (1 - \theta)^{-1}/\gamma$. \square

Corollary 2 (of Theorem 5) *Suppose the Oracle Mechanism has access to a γ -approximation oracle. Then, if instead of budget B , the mechanism is given a reduced budget γB , it would be an almost budget feasible mechanism with approximation ratio $\gamma^2/2$.*

E.4 An Improved Polynomial-Time Mechanism

In this section, we present a polynomial-time mechanism with approximation ratio $\frac{1}{3}$. The mechanism follows the idea of the oracle mechanism, except that instead of computing the optimum cost per utility rate (which requires exponential running time), the mechanism computes an estimate for the cost per utility rate, which we call the *stopping rate*. Before presenting the mechanism, we need to formally define the notion of stopping rate.

Definition 13 *Suppose we are given an instance of the problem with cost vector c , submodular utility function F and budget B . Construct the sequence $\chi(F)$ and choose the largest integer k such that $F(\chi_k) \cdot \frac{c_{x_k}}{\partial_k} \leq B$. The stopping rate, denoted by $r(c, F, B)$, is then defined to be $B/F(\chi_k)$. We sometimes denote the stopping rate simply by $r(B)$ when c, F are clear from the context.*

Now we define the mechanism by presenting the winner selection and payment rules.

Winner Selection The Mechanism constructs the sequence $\chi(F)$ and chooses the largest integer k such that $F(\chi_k) \cdot \frac{c_{x_k}}{\partial_k} \leq B/2$. Then, it reports χ_k as the set of winners.

The Payment Rule For simplicity, assume that the winners are indexed from 1 to k . The Payment to winner i is equal to $2r_i \cdot \partial_i$, where $r_i = r(c', F, B/2)$ and c' is the cost vector which is identical to the original cost vector c except that c'_i is set to ∞ (i.e. intuitively, seller i is removed from S). In other words, think of an auxiliary instance in which we are given budget $B/2$ and seller i is removed from the set of sellers S . Then, r_i is the stopping rate in this instance.

The above definitions are formally put together in Mechanism 2. We also address our polynomial-time mechanism as Mechanism 2. In the rest of this section, we prove that Mechanism 2 is individually rational, oneway-truthful and budget feasible, and also, it has approximation ratio $1/3$. As we mentioned before, this mechanism can be converted to a (fully) truthful and strictly budget feasible mechanism without any (asymptotic) loss in the approximation ratio; the details are discussed in Sections E.5 and E.6.

Mechanism 2: A polynomial-time $(1/3)$ -approximation mechanism

input : Submodular utility function F , Budget B

Sellers report their costs;
Construct the sequence $\chi(F)$;
Find the smallest integer k such that $F(\chi_k) \cdot \frac{c_{x_k}}{\partial_k} \leq B/2$;
Announce χ_k as the set of winners;
foreach $i \in \chi_k$ **do**
 $c' \leftarrow c$;
 $c'_i \leftarrow \infty$;
 $r_i \leftarrow r(c', F, B/2)$;
 Pay $r_i \cdot \partial_i$ to seller i ;
end

Simplifying Assumption Through out the analysis, w.l.o.g. we assume that the sellers appear in the greedy sequence in an increasing order, i.e. $x_i = i$ for all $i \in S$.

Lemma 19 *Mechanism 2 is individually rational.*

Proof. According to the payment rule, it is enough to show that for each winner i we have $r_i \geq c_i/\partial_i$. Let $\bar{\chi}(F) = \langle \bar{x}_1, \dots, \bar{x}_{n-1} \rangle$ denote the greedy sequence for when i is removed. Also, let $\bar{\chi}_j$ denote the subset containing first j elements of $\bar{\chi}(F)$, and $\bar{\partial}_j = F(\bar{\chi}_j) - F(\bar{\chi}_{j-1})$.

Note that $x_j = \bar{x}_j$ for all $j < i$. Now, consider the following two cases for the proof: Let \bar{k} be the largest integer j satisfying $F(\bar{\chi}_j) \cdot c_{\bar{x}_j}/\bar{\partial}_j \leq B/2$. We either have that $\bar{k} = i - 1$ or $\bar{k} \geq i$. We prove the lemma by proving that $r_i \geq c_i/\partial_i$ in each of these cases.

First, suppose $\bar{k} = i - 1$. This means $r_i = \frac{B}{2F(\chi_{i-1})}$. Also, note that $F(\chi_i) \cdot \frac{c_i}{\partial_i} \leq B/2$ (since $i < k$). The two latter facts, along with the fact that $F(\chi_{i-1}) \leq F(\chi_i)$ together imply that $r_i \geq c_i/\partial_i$.

It remains to do the proof in the second case: Suppose $\bar{k} \geq i$. Note that due to the definition of the greedy sequence we have $c_{x_i}/\partial_i \leq c_{\bar{x}_i}/\bar{\partial}_i$. So, the proof is done if we show that $r_i \geq c_{\bar{x}_i}/\bar{\partial}_i$. To this end, first observe that we have $F(\bar{\chi}_i) \cdot c_{\bar{x}_i}/\bar{\partial}_i \leq B/2$ due to the fact that $\bar{k} \geq i$. This just implies $c_{\bar{x}_i}/\bar{\partial}_i \leq B/(2F(\bar{\chi}_{\bar{k}}))$ due to the monotonicity of F . Finally, seeing that $B/(2F(\bar{\chi}_{\bar{k}})) = r_i$ proves the claim. \square

Lemma 20 *Mechanism 2 is oneway-truthful.*

Proof. The proof is almost identical to the proof of Lemma 16, we state the proof for completeness. For contradiction, suppose there exists a seller $i \in S$ who has incentive to report a cost \bar{c}_i which is higher than her true cost c_i . Assuming that the mechanism picked χ_k as the set of winners, we have either $i > k$ or $i \leq k$. The proof is done separately in each of these cases.

If $i > k$, then see that seller i can not change the first $i - 1$ elements of $\chi(F)$ by reporting a higher cost. Now, see that if $i > k + 1$, then again χ_k will be chosen as the set of winners, which is

a contradiction with the incentive of seller i for misreporting. If $i = k + 1$, then see that χ_k will be chosen as a subset of the winners; in this case, the set of winners can possibly contain other sellers, but not certainly not seller $k + 1$ (due to the fact that she was not chosen before). Consequently, seller i remains a loser even by reporting \bar{c}_i . Contradiction.

It remains to do the proof for when $i \leq k$. Recall that the payment to seller i is $r_i \cdot \partial_i$. Since r_i is not a function of the cost reported by seller i , then see that the only way that i can increase her utility by misreporting, is increasing ∂_i . But by reporting a cost higher than c_i , seller i can not change the first $i - 1$ elements of $\chi(F)$, and so, can only decrease ∂_i . Consequently, the payment to i does not increase in this case as well which means i has no incentive to report a higher cost. \square

Lemma 21 *Mechanism 2 has approximation ratio $1/3$ in large markets.*

Proof. First we prove that $F(\chi_{k+1}) \geq F^*/3$. This would imply $F(\chi_k) \geq F^* \cdot (1/3 - \theta)$ due to the large market assumption, which proves the lemma.

So, it is enough to show that $F(\chi_{k+1}) \geq F^*/3$. We prove this by contradiction, assume $F(\chi_{k+1}) < F^*/3$. For more intuition, we first explain our argument for contradiction in words and then we state it more formally. From the definition of the greedy sequence $\chi(F)$, it can be seen that conditioned on buying the subset χ_{k+1} , the cost for buying each extra unit of utility is at least $\frac{c_{k+1}}{\partial_{k+1}}$. So, even if we get the subset χ_{k+1} for free, the cost for buying an additional $\frac{2F^*}{3}$ units of utility (which is needed for the optimum solution) would be at least

$$\frac{2F^*}{3} \cdot \frac{c_{k+1}}{\partial_{k+1}} > 2F(\chi_{k+1}) \cdot \frac{c_{k+1}}{\partial_{k+1}} \geq B.$$

This means cost of the optimum solution is more than B .

To formalize this contradiction, just note that by monotonicity of F , we have

$$F(S^* \cup \chi_{k+1}) - F(\chi_{k+1}) > \frac{2F^*}{3}. \quad (28)$$

Now observe that by the definition of the greedy sequence $\chi(F)$, the cost for buying each extra unit of utility conditioned on having χ_{k+1} is at least $\frac{c_{k+1}}{\partial_{k+1}}$. This fact, and (28) together imply that

$$c(S^* \cup \chi_{k+1}) - c(\chi_{k+1}) > \frac{c_{k+1}}{\partial_{k+1}} \cdot \frac{2F^*}{3}. \quad (29)$$

On the other hand, recall that $F(\chi_{k+1}) < F^*/3$, so we can write (29) as

$$\begin{aligned} c(S^* \cup \chi_{k+1}) - c(\chi_{k+1}) &> \frac{c_{k+1}}{\partial_{k+1}} \cdot \frac{2F^*}{3} \\ &> 2F(\chi_{k+1}) \cdot \frac{c_{k+1}}{\partial_{k+1}} > B. \end{aligned}$$

This implies $c(S^*) > B$ which is a contradiction with the budget feasibility of S^* . \square

We prove that Mechanism 2 is almost budget feasible in Lemma 23. Before that, we first need to prove the following lemma which will be used in the proof of Lemma 23. This Lemma states that $F(\chi_k)$ is (roughly) at least half of the optimal utility in a large market.

Lemma 22 *Suppose we are given a θ -large market with $\theta \leq 1/2$ and let χ_k denote the subset chosen by Mechanism 2 when run on this instance. Then we have*

$$F(\chi_k) \geq (1/2 - \theta) \cdot F^*$$

Proof. Recall that w.l.o.g. we assume $\chi = \langle 1, \dots, n \rangle$. To prove the lemma, it is enough to show that $F(\chi_{k+1}) \geq F^*/2$.

$$F(\chi_k) \geq F(\chi_{k+1}) - \theta \cdot F^* \geq (1/2 - \theta) \cdot F^*,$$

where the first inequality is due to the large market assumption.

We prove $F(\chi_{k+1}) \geq F^*/2$ by contradiction, assume $F(\chi_{k+1}) < F^*/2$. For more intuition, we first explain our argument for contradiction in words and then we state it more formally. Conditioned on buying the subset χ_{k+1} , the cost for buying each extra unit of utility is at least $\frac{c_{k+1}}{\partial_{k+1}}$. So, even if we get the subset χ_{k+1} for free, the cost for buying an additional $\frac{F^*}{2}$ units of utility (which is needed for the optimum solution) would be strictly more than $\frac{F^*}{2} \cdot \frac{c_{k+1}}{\partial_{k+1}} \geq B$. To formalize this contradiction, just note that by monotonicity of F , we have

$$F(S^* \cup \chi_{k+1}) - F(\chi_{k+1}) > \frac{F^*}{2}. \quad (30)$$

Now observe that by the definition of the greedy sequence $\chi(F)$, the cost for buying each extra unit of utility conditioned on having χ_{k+1} is at least $\frac{c_{k+1}}{\partial_{k+1}}$. This fact, and (30) together imply that

$$c(S^* \cup \chi_{k+1}) - c(\chi_{k+1}) > \frac{c_{k+1}}{\partial_{k+1}} \cdot \frac{F^*}{2} \geq B.$$

This implies $c(S^*) > B$ which is a contradiction with the budget feasibility of S^* . \square

Lemma 23 *Mechanism 2 is almost budget feasible*

Proof. Let r denote $r(c, F, B/2)$, which is equal to $\frac{B/2}{F(\chi_k)}$ by definition. We show that for any $i \in \chi_k$, we have $r_i \leq 2r \cdot (1 - 3\theta)^{-1}$. This would prove that sum of the payments is at most $B \cdot (1 - 3\theta)^{-1}$

since:

$$\sum_{i=1}^k r_i \partial_i \leq (1 - 3\theta)^{-1} \cdot \sum_{i=1}^k 2r \partial_i = (1 - 3\theta)^{-1} \cdot 2r F(\chi_k) = (1 - 3\theta)^{-1} \cdot B.$$

This would finish the proof due to the large market assumption.

To this end, fix a seller $i \in S$ and let $\bar{\chi}(F) = \langle \bar{x}_1, \dots, \bar{x}_{n-1} \rangle$ denote the greedy sequence for when i is removed from the set of sellers S . Also, let $\bar{\chi}_j$ denote the subset containing first j elements of $\bar{\chi}(F)$, and $\bar{\partial}_j = F(\bar{\chi}_j) - F(\bar{\chi}_{j-1})$. Finally, let \bar{k} be the largest integer j satisfying $F(\bar{\chi}_j) \cdot c_{\bar{x}_j} / \bar{\partial}_j \leq B/2$.

Suppose $F_{B/2}^*$ denotes the optimal utility for when the budget is reduced to $B/2$. Then we would have:

$$r_i = \frac{B/2}{F(\bar{\chi}_{\bar{k}})} \leq \frac{B/2}{(1/2 - \theta) \cdot F_i^*} \tag{31}$$

$$\leq \frac{B/2}{(1/2 - \theta) \cdot (1 - \theta) \cdot F_{B/2}^*} \tag{32}$$

$$\leq \frac{B/2}{(1/2 - \theta) \cdot (1 - \theta) \cdot F(\chi_k)}$$

$$\leq \frac{r}{(1/2 - \theta) \cdot (1 - \theta)}$$

$$\leq \frac{2r}{1 - 3\theta}$$

where (31) is due to Lemma 22 and (32) is due to the θ -largeness of the market. \square

E.5 From Oneway-Truthfulness to Truthfulness

Changing the payment rule is the key to convert all of our oneway-truthful mechanisms to truthful mechanisms. In fact, the winner selection rule remains identically the same, however, we replace the payment rule in all of our mechanisms by the following simple payment rule:

The Critical Cost Payment Rule Each winner i is paid the highest cost that it could report and still remain a winner.

Individual rationality and truthfulness of all the mechanisms are trivial under this payment rule. Also, the proofs for approximation ratio remain the same since we have not changed the winner selection rule. The only non-trivial part is proving the budget feasibility under this new payment rule. For all of the mechanisms in this paper, we can show that replacing the original payment rule with the Critical Cost payment rule, does not increase the payment of the mechanism to any of the winners. First we prove this for the Oracle mechanism.

Lemma 24 *In the Oracle mechanism, the payment to each winner does not increase if we replace the payment rule with the Critical Cost payment rule.*

Proof. We need to show that the Critical Cost payment rule does not pay to winner i more than $2r_i \cdot \partial_i$. Equivalently, we show that if i reports a cost $\bar{c}_i > 2r_i \cdot \partial_i$, then it will not be selected as a winner anymore. The proof is by contradiction, suppose i reports a cost higher than $\bar{c}_i > 2r_i \cdot \partial_i$ and still gets selected as a winner.

Let the instance in which the cost c_i is replaced by \bar{c}_i be called the *fake* instance. Let \bar{c} denote the cost vector in the fake instance which is identical to the original cost vector c except that c'_i is set to \bar{c}_i . For any subset of sellers $S' \subseteq S$, let $\bar{c}(S') = \sum_{s \in S'} \bar{c}_s$.

Also, let $\bar{\chi}(F)$ denote the greedy sequence for the fake instance, $\bar{\chi}_j$ denote the subset containing first j elements of $\bar{\chi}(F)$, and $\bar{\partial}_j = F(\bar{\chi}_j) - F(\bar{\chi}_{j-1})$. Finally, let \bar{S}^* denote the optimum budget feasible subset in the fake instance.

For more intuition, we first explain our argument for contradiction in words and then we state it more formally. See that the optimum utility is at least F_i^* in the fake instance (recall that F_i^* denotes the optimum utility when i is removed from S), which implies $F(\bar{S}^*) - F(\bar{\chi}_{\bar{k}}) \geq \frac{F_i^*}{2}$. Also, we will show that even if we get $\bar{\chi}_{\bar{k}}$ for free, the cost for buying an additional $F(\bar{S}^*) - F(\bar{\chi}_{\bar{k}})$ units of utility would be more than $2r_i$ per unit of utility. The two latter facts together would imply that

$$\bar{c}(\bar{S}^* \setminus \bar{\chi}_{\bar{k}}) > 2r_i \cdot (F(\bar{S}^*) - F(\bar{\chi}_{\bar{k}})) \geq 2r_i \cdot \frac{F_i^*}{2} = B$$

which implies $c(\bar{S}^*) > B$. This is a contradiction with budget feasibility of \bar{S}^* .

To formalize this contradiction, just note that by the monotonicity of F , we have

$$F(\bar{S}^* \cup \bar{\chi}_{\bar{k}}) - F(\bar{\chi}_{\bar{k}}) \geq \frac{F_i^*}{2}. \quad (33)$$

Now observe that by the definition of the greedy sequence $\bar{\chi}(F)$, the cost for buying each extra unit of utility, conditioned on having the subset $\bar{\chi}_{\bar{k}}$, is at least $\bar{c}_i / \bar{\partial}_{\bar{k}}$, i.e. the cost of the seller at position \bar{k} of $\bar{\chi}(F)$ divided by the marginal utility that it adds. Also, see that $\bar{c}_i / \bar{\partial}_{\bar{k}} > 2r_i$ since $\bar{c}_i > 2r_i \cdot \partial_i$ and $\bar{\partial}_{\bar{k}} \leq \partial_i$. This proves that the cost for buying each extra unit of utility, conditioned on having the subset $\bar{\chi}_{\bar{k}}$, is more than $2r_i$. This fact, and (33) together imply that

$$\bar{c}(\bar{S}^* \cup \bar{\chi}_{\bar{k}}) - \bar{c}(\bar{\chi}_{\bar{k}}) > 2r_i \cdot \frac{F_i^*}{2} = B.$$

This implies $\bar{c}(\bar{S}^*) > B$ which is a contradiction with the budget feasibility of \bar{S}^* . \square

Now, we prove an equivalent version of Lemma 24 for Mechanism 2.

Lemma 25 *In Mechanism 2, the payment to each winner does not increase if we replace the payment rule with the Critical Cost payment rule.*

Proof. We need to show that the Critical Cost payment rule does not pay to winner i more than $r_i \cdot \partial_i$. Equivalently, we show that if i reports a cost x with $x > r_i \cdot \partial_i$, then it will not be selected as a winner anymore. The proof is by contradiction, suppose i reports a cost x which is higher than $r_i \cdot \partial_i$ and still gets selected as a winner.

Let the instance in which the cost c_i is replaced by x be called the *fake* instance. Let \bar{c} denote the cost vector in the fake instance which is identical to the original cost vector c except that \bar{c}_i is set to x . For any subset of sellers $S' \subseteq S$, let $\bar{c}(S') = \sum_{s \in S'} \bar{c}_s$; when there is no risk of confusion, we sometimes denote $c(\{i\})$ by $c(i)$ for any cost function $c(\cdot)$.

Also, consider the instance in which the cost c_i is replaced by ∞ (seller i is removed) and call it the *large* instance. Let $\bar{\bar{c}}$ denote the cost vector in the large instance which is identical to the original cost vector c except that $\bar{\bar{c}}_i$ is set to ∞ . For any subset of sellers $S' \subseteq S$, let $\bar{\bar{c}}(S') = \sum_{s \in S'} \bar{\bar{c}}_s$.

Let $\bar{\chi}(F) = \langle \bar{x}_1, \dots, \bar{x}_n \rangle$ and $\bar{\bar{\chi}}(F) = \langle \bar{\bar{x}}_1, \dots, \bar{\bar{x}}_{n-1} \rangle$ respectively denote the greedy sequences for the fake and large instances, and $\bar{\chi}_j, \bar{\bar{\chi}}_j$ respectively denote the subsets containing the first j elements of $\bar{\chi}(F)$ and $\bar{\bar{\chi}}(F)$. Also, let $\bar{\partial}_j = F(\bar{\chi}_j) - F(\bar{\chi}_{j-1})$ and $\bar{\bar{\partial}}_j = F(\bar{\bar{\chi}}_j) - F(\bar{\bar{\chi}}_{j-1})$.

Suppose that Mechanism 2 picks $\bar{\chi}_{\bar{k}}$ as the set of winners in the fake instance and $\bar{\bar{\chi}}_{\bar{k}}$ as the set of winners in the large instance, i.e. the mechanism picks the first \bar{k} elements of $\bar{\chi}(F)$ and the first \bar{k} elements of $\bar{\bar{\chi}}(F)$ as the set of winners in each of the instances. First, we need to prove the following claim and then we proceed to the proof of the lemma.

Claim 3 *The first \bar{k} elements of $\bar{\chi}(F)$ and $\bar{\bar{\chi}}(F)$ are identical.*

Proof Let \bar{i} denote the position of seller i in $\bar{\chi}(F)$, i.e. we have $\bar{x}_{\bar{i}} = i$. Then, see that

$$c(\bar{x}_{\bar{k}})/\bar{\partial}_{\bar{k}} \leq r_i, \quad (34)$$

which holds since we have $F(\bar{\chi}_{\bar{k}}) \cdot c(\bar{x}_{\bar{k}})/\bar{\partial}_{\bar{k}} \leq B/2$ by the definition of \bar{k} and $r_i = B/(2F(\bar{\chi}_{\bar{k}}))$ by the definition of stopping rate.

On the other hand, see that

$$\begin{aligned} \bar{c}(x_{\bar{i}})/\bar{\partial}_{\bar{i}} &> r_i \cdot \partial_i/\bar{\partial}_{\bar{i}} \\ &\geq r_i \end{aligned} \quad (35)$$

where (35) is due to the fact that $\partial_i \geq \bar{\partial}_{\bar{i}}$; this fact holds because the first $i - 1$ elements of $\chi(F)$ and $\bar{\chi}(F)$ are identical, which means the marginal utility added by seller i in $\chi(F)$ is more than her marginal in $\bar{\chi}(F)$.

To summarize, see that by (34) and (35) we have:

$$\begin{aligned} c(\bar{x}_{\bar{k}})/\bar{\partial}_{\bar{k}} &\leq r_i, \\ \bar{c}(\bar{x}_{\bar{i}})/\bar{\partial}_{\bar{i}} &> r_i \end{aligned}$$

which means seller i will be never be used in its first \bar{k} positions of $\bar{\chi}(F)$ (due to the greedy construction of the sequence). Consequently, the first \bar{k} elements of $\bar{\chi}(F)$ and $\bar{\bar{\chi}}(F)$ are identical. ■

Now we are ready to see to the contradiction, it follows from the following set of inequalities which will be clarified below.

$$B/2 \geq F(\bar{\chi}_{\bar{k}}) \cdot \bar{c}_{\bar{k}}/\bar{\partial}_{\bar{k}} \tag{36}$$

$$> F(\bar{\chi}_{\bar{k}}) \cdot \frac{r_i \cdot \partial_i}{\bar{\partial}_{\bar{i}}} \tag{37}$$

$$\geq F(\bar{\chi}_{\bar{k}}) \cdot r_i \tag{38}$$

$$\geq F(\bar{\bar{\chi}}_{\bar{k}}) \cdot r_i \tag{39}$$

$$=B/2$$

(36) is due to the definition of \bar{k} in the fake instance;

(37) holds since we have assumed $\bar{i} \leq \bar{k}$ (i.e. i wins in the fake instance) and so we have $\bar{c}(\bar{x}_{\bar{i}})/\bar{\partial}_{\bar{i}} \leq \bar{c}_{\bar{k}}/\bar{\partial}_{\bar{k}}$ by the definition of the greedy sequence, this fact, and the fact that $\bar{c}(\bar{x}_{\bar{i}}) > r_i \cdot \partial_i$ directly imply (37);

(38) holds since $\partial_i \geq \bar{\partial}_{\bar{i}}$; we already proved this is true in the proof of Claim 3: because the first $i - 1$ elements of $\chi(F)$ and $\bar{\chi}(F)$ are identical;

(39) follows by the monotonicity of F and since $\bar{\bar{\chi}}_{\bar{k}} \subseteq \bar{\chi}_{\bar{k}}$ (which holds by Claim 3). □

E.6 Strictly Budget Feasible Mechanisms

In this section, we show how to convert our almost budget feasible mechanisms to strictly budget feasible mechanisms. We start with the Oracle Mechanism, i.e. Mechanism 1. Recall that in Lemma 18, we proved that sum of the payments in the Oracle mechanism is at most $B \cdot (1 - \theta)^{-1}$. So, if instead of budget B , we give a reduced budget $B \cdot (1 - \theta)$ to the oracle mechanism, then the sum of its payments would not exceed B .

It only remains to show that this budget reduction does not affect the approximation ratio significantly. To this end, first we prove that the budget reduction does not affect the optimum utility significantly.

Lemma 26 *Assume we are given a θ -large market and Let F_b^* denote the optimal solution for*

this market when the budget is reduced to $b \leq B$. Then for any given constant $\epsilon > 0$ and $b = B \cdot (1 - \epsilon)$ we have:

$$F_b^* \geq (1 - \theta - \epsilon) \cdot F^*$$

Proof. W.L.O.G. suppose that with budget B , the optimal subset is $S^* = \{1, \dots, s\}$. Let the submodular function G denote the restriction of F to the subset S^* , i.e. $G = F|_{S^*}$. Now, construct the greedy sequence $\chi(G)$, and denote it by $\chi(G) = \langle 1, \dots, s \rangle$, where w.l.o.g. we have assumed that the members of S^* appear in the greedy sequence in the increasing order.

By assumption, we have that $c(S^*) \leq B$. Let s' denote the smallest integer such that

$$\sum_{i=1}^{s'} c_i \geq B \cdot (1 - \epsilon).$$

Also, let $S' = \{1, \dots, s' - 1\}$. We claim that

$$F(S') \geq (1 - \theta - \epsilon) \cdot F^*,$$

which proves the lemma since $c(S') \leq B \cdot (1 - \epsilon)$. To this end, first verify that

$$F(S' \cup \{s'\}) \geq F^* \cdot (1 - \epsilon);$$

this holds by the definition of the greedy sequence, which picks the cheaper items (relative to their utility) first. Then, observe that $F(S') \geq F(S' \cup \{s'\}) - F(\{s'\})$ due to submodularity. This implies

$$F(S') \geq (1 - \epsilon) \cdot F^* - \theta \cdot F^* = F^* \cdot (1 - \theta - \epsilon)$$

due to the large market assumption. □

To achieve a strictly budget feasible mechanism for large markets, fix ϵ to be an arbitrary small constant. In a large market with $\theta < \epsilon$, we can reduce the budget of the Oracle mechanism to $B \cdot (1 - \epsilon)$ to get a strictly budget feasible mechanism: the Oracle Mechanism would be budget feasible by Lemma 18. Also, by Lemma 26, such a market (with the reduced budget) would be $\frac{\theta}{1 - \theta - \epsilon}$ -large. So, by Lemmas 17 and 26, the approximation ratio of the Oracle mechanism would be

$$\left(\frac{1}{2} - \frac{\theta}{1 - \theta - \epsilon} \right) \cdot (1 - \theta - \epsilon).$$

See that in large markets, where $\theta \rightarrow 0$, the approximation ratio would be

$$\lim_{\theta \rightarrow 0} \left(\frac{1}{2} - \frac{\theta}{1 - \theta - \epsilon} \right) \cdot (1 - \theta - \epsilon) = \frac{1}{2} \cdot (1 - \epsilon),$$

i.e. for any arbitrary small $\epsilon > 0$, we have a large market mechanism with approximation ratio $\frac{1}{2} \cdot (1 - \epsilon)$.

By a similar argument, it can be seen that for any arbitrary small constant $\epsilon > 0$, Mechanism 2 can also be converted to a strictly budget feasible mechanism with approximation ratio $\frac{1}{3} \cdot (1 - \epsilon)$ in large markets.

F Hoeffding Bounds

In this section we state a version of Hoeffding bounds (HFFDNGBND) that is suitable for our purpose.

Hoeffding Bounds. Let x_1, \dots, x_n be i.i.d. random variables such that $\Pr(x_i \in [a, b]) = 1$. Let $\mu = \mathbb{E}[\sum_{i=1}^n x_i]$, then we have:

$$\Pr\left(\sum_{i=1}^n x_i \geq (1 + \epsilon) \cdot \mu\right) \leq e^{-\frac{2\epsilon^2\mu^2}{n(b-a)^2}}.$$