Budget Feasible Procurement Auctions

Nima Anari
University of California, Berkeley, anari.nima@gmail.com

Gagan Goel
gagangoel@gmail.com

Afshin Nikzad
Stanford University, nikzad@stanford.edu

We consider a simple and well-studied model for procurement problems and solve it to optimality. A buyer with a fixed budget wants to procure, from a set of available workers, a budget feasible subset that maximizes her utility: Any worker has a private reservation price and provides a publicly known utility to the buyer in case of being procured. The buyer's utility function is additive over items. The goal is designing a direct revelation mechanism that solicits workers' reservation prices and decides which workers to recruit and how much to pay them. Moreover, the mechanism has to maximize the buyer's utility without violating her budget constraint. We study this problem in the prior-free setting; our main contribution is finding the optimal mechanism in this setting, under the "Small Bidders" assumption. This assumption, also known as the "small bid to budget ratio assumption", states that the bid of each seller is small compared to the buyer's budget. We also study a more general class of utility functions: submodular utility functions. For this class, we improve the existing mechanisms significantly under our assumption.

Key words: auction design, robust auctions, procurement, budget feasible, knapsack problem

Subject classification: Games/group decisions: Bidding/auctions

Contents

1. Introduction

We consider a simple and well-studied model for procurement problems and solve it to optimality. A buyer with a fixed budget wants to procure, from a set of available workers, a "budget feasible subset" that maximizes her utility: A worker has a private reservation price and provides a publicly known utility to the buyer in case of being procured. The goal is designing a direct revelation mechanism that solicits workers' reservation prices and decides which workers to recruit and how much to pay them. Moreover, the mechanism has to maximize the buyer's utility without violating her budget constraint.

With complete information, the buyer faces the so-called *knapsack* problem: If all the reservation prices were known to the buyer, then she could choose, among all the "budget feasible subsets" of workers, the one that maximizes her utility (where a budget feasible subset is a subset for which the sum of reservation prices of the workers belonging to it does not exceed the buyer's budget). Since the reservation prices are private information in many applications, some of which are discussed below, we take the mechanism design approach to study this problem.

This basic model has been used for various applications, namely, in allocating R&D subsidies by government agencies Ensthaler and Giebe (2014b), conducting auctions for reducing the emission of greenhouse gases Maskin (2002), pricing tasks in social networks Singer and Mittal (2013), experiment design Horel et al. (2013), and marketing over social networks Singer (2011). (Section 10 contains a detailed discussion of these applications.)

If the exact distributions of the sellers' reservation prices are given, it is possible to find optimal *ex ante* budget feasible mechanisms (Ensthaler and Giebe (2014a)). This, however, is not a reasonable assumption for our applications, as it is not always possible to find reasonably accurate distributions (see the discussion of the applications in Section 10). We therefore look for the optimal mechanism in the prior-free setting, which would be robust such uncertainties. Even though this model has a wide range of applications, it has never been solved to optimality, to the extent of our knowledge.

We measure the performance of a mechanism by its *competitive ratio*, which is the fraction of the utility of the underlying knapsack problem that the mechanism is always guaranteed to attain. Our main contribution is finding the mechanism that attains the highest competitive ratio among all the truthful mechanisms, under the "Small Bidders" assumption. Informally speaking, this assumption means that the reservation price of each worker is small compared with the overall budget. We will see that this is a reasonable assumption in all of the mentioned applications (Section 10); thus, our results apply to these applications as well.

2. Setup

We define the model abstractly. Consider a reverse auction 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, with a corresponding cost c_i that is only known to herself. We use $\mathbf{c} \in (\mathbb{R}_+)^n$ to denote $(c_1, \ldots, c_n)$. The buyer derives a publicly known utility u_i from item i, and her utility function is additive over the items, i.e. her utility from owning a subset of items is the sum of the utilities of the items in that subset. The buyer has a limited budget B, and her goal is to buy a subset of items that maximizes her total utility without exceeding her budget. Note that the spent budget is not part of the buyer's objective function, it is only a constraint.

The underlying optimization problem, in which the sellers are not strategic and the costs are known to the buyer, is the well-known *knapsack* problem.[1] Unlike the knapsack problem, we assume that the cost c_i is private information of seller i. We are interested in designing direct-revelation mechanisms in which the buyer solicits bids from the sellers, and then computes which sellers to buy from and how much to pay them. We consider a prior-free setting, i.e. we suppose the buyer has no prior information about the costs (unlike the Bayesian setting).

More formally, a mechanism $\mathcal{M}$ consists of two functions $A : (\mathbb{R}_+)^n \to \{0,1\}^n$ and $P : (\mathbb{R}_+)^n \to (\mathbb{R}_+)^n$. The *selection rule* $A(\cdot)$ takes as input the costs of n sellers and reports the set of winners (sellers whose item will be purchased). The payment rule $P(\cdot)$ takes as input the costs of n sellers and reports how much is paid to each of them. We will use functions $A_i : (\mathbb{R}_+)^n \to \{0,1\}$ and $P_i : (\mathbb{R}_+)^n \to \mathbb{R}_+$, 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, and are called the selection rule and payment rule for seller i, respectively.

We require any mechanism $\mathcal{M} = (A, P)$ to 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 $\mathbf{c} \in (\mathbb{R}_+)^n$.
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 for sellers, i.e. for all non-truthful reports $\overline{c}_i$ from seller i, it holds that

$$P_i(\overline{c}_i, \mathbf{c_{-i}}) - c_i \cdot A_i(\overline{c}_i, \mathbf{c_{-i}}) \leq P_i(c_i, \mathbf{c_{-i}}) - c_i \cdot A_i(c_i, \mathbf{c_{-i}}) \tag{1}$$

Defining a Benchmark. Among all mechanisms that satisfy the above properties, we are interested to the one that maximizes the utility of the buyer with respect to the following benchmark. Let $U^*(\mathbf{c}, \mathbf{u})$ denote the utility of the omniscient mechanism, i.e. the utility of the knapsack optimization problem assuming that costs of the sellers are known to the buyer.[2] When there is no risk of confusion, we also denote $U^*(\mathbf{c}, \mathbf{u})$ by U^* for brevity.

Definition 1. A mechanism $\mathcal{M}$ is α-*competitive* when $\alpha \in [0, 1]$ is the largest scalar for which the mechanism derives utility at least $\alpha \cdot U^*(\mathbf{c}, \mathbf{u})$ for all $\mathbf{c}$ and $\mathbf{u}$.

Our main contribution is finding the mechanism that attains the highest possible competitive ratio in the class of truthful mechanisms.

2.1. The Small Bidders Assumption

Our small bidders assumption states that the cost of a single item is very small compared to the buyer's budget B.

The Small Bidders Assumption. Assume that $c_{\mathsf{max}} \ll B$, where $c_{\mathsf{max}} = \max_{i \in S}\{c_i\}$.

An alternative way to write the assumption is $c_{\mathsf{max}} = o(B)$; in other words, we define *bid-budget ratio* of the market to be $\theta = \frac{c_{\mathsf{max}}}{B}$ and analyze our mechanisms for when $\theta \to 0$. Our mechanisms, however, do not need "very small" θ to perform well; this is elaborated during the discussion of our results in Section 3, where we note that even for θ as large as $1/20$ our mechanisms have a very close performance to the optimum performance.

This assumption is also known as the *small bid to budget ratio assumption* and is used in other 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, i.e. mechanisms for additive utility functions, will be analyzed under this assumption. Throughout the paper we assume that the competitive ratio of a mechanism denotes its (asymptotic) competitive ratio under this assumption.

2.2. Extension: Divisible Items.

In this section, we extend our model for indivisible items to the case of divisible items. There are two reasons for studying the case of divisible items. First, our mechanisms for indivisible items are based on our mechanisms for divisible items. Second, the case of divisible items is interesting by itself; for instance, if the item being sold by a seller is her own time (e.g. consider crowdsourcing markets) then it can be modeled as a divisible item.

We verify that our model for indivisible items can be naturally extended to divisible items. The cost of a fraction $x \leq 1$ of the item of seller i is $x \cdot c_i$ to her, and the utility that the buyer attains from that fraction is $x \cdot u_i$. Range of the selection rule is $[0,1]^n$, i.e. the selection rule is defined as $A : (\mathbb{R}_+)^n \to [0,1]^n$. Individual rationality states that seller i should be paid at least $x \cdot c_i$ in exchange for a fraction x of her item. Truthfulness is defined by (1).

2.3. Extension: Submodular Utility Functions.

A well-studied generalization is when the buyer's utility function over the set of items is a submodular function rather than an additive function. More precisely, if we denote the utility function by $U : 2^S \to \mathbb{R}_+$, then for additive utility functions we have $U(T) = \sum_{i \in T} u_i$, for $\forall T \subseteq S$. We also study a more general case when $U(\cdot)$ is submodular. Our mechanisms for submodular utility functions work under a different Small Bidders assumption which we state below.

The Alternative Small Bidders Assumption. Assume that $u_{\mathsf{max}} \ll U^*$, where $u_{\mathsf{max}} = \max_{i \in S} u_i$.

Recall that U^* is the optimum utility, i.e. utility of the omniscient mechanism. An alternative way to write the assumption is $u_{\mathsf{max}} = o(U^*)$; in other words, we define *largeness ratio* of the market to be $\theta = \max_{i \in S} \frac{u_i}{U^*}$ and analyze our mechanisms for when $\theta \to 0$.

It is worth pointing out that we can slightly modify our mechanisms for additive utility functions so that they work under this alternative assumption as well, while preserving their competitive ratio. Also, we note that our impossibility result for additive utilities (Section 7) holds for either of the assumptions.

3. Organization

We design optimal budget-feasible mechanisms under the Small Bidders assumption. To the best of our knowledge, this work is the first to study the case of small bidders. We list our results below. All of the mechanisms we provide satisfy the basic properties: budget feasibility, individual rationality, and truthfulness.

1. For divisible items, we design a deterministic mechanism with competitive ratio $1 - 1/e$ (Section 5). Previously, no mechanism was known for the case of divisible items. In fact, one can show that all truthful mechanisms for divisible items would have competitive ratio 0 if the Small Bidders assumption is dismissed.
2. For indivisible items, we provide a randomized truthful mechanism with competitive ratio $1 - 1/e$. (Section 8)
3. We show the optimality of the above mechanisms by proving that no truthful (and possibly randomized) mechanism can achieve competitive ratio better than $1 - 1/e$. Our impossibility result holds for both cases of divisible and indivisible items. (Section 7)
4. For the case of submodular utility functions, we design deterministic mechanisms that achieve competitive ratios $\frac{1}{2}$ and $\frac{1}{3}$ with exponential and polynomial running times respectively. We only consider the case of indivisible items. (Section 9)

Our optimal mechanism differs from most of the previous mechanisms proposed in the literature in that, rather than using *cutoff allocation rules* and their corresponding *linear payment rules*, it uses a non-linear payment rule corresponding to a non-cutoff allocation rule (see Sections 4 and 8 for detailed comparisons). Our optimal mechanism for indivisible items uses our mechanism for divisible items to find a *fractional allocation*, and then implements it as a lottery over integral allocations, while guaranteeing incentive compatibility, ex post budget feasibility, and ex post individual rationality.

It is also worth pointing out that the impossibility result for additive utility functions also holds for submodular utility functions, and it is the best impossibility result known to us for submodular utility functions. Therefore, unlike the case of additive utility functions, we do not provide matching upper and lower bounds for the case of submodular utility functions.

Recall from Section 2.1, the definition of bid-budget ratio, denoted by θ. In our analysis, and also in the applications that we focus on (Section 10), we consider θ to be "small". For analytical tractability, we focus on the case of $\theta \to 0$ and state our main theorems for this setting. However, our mechanisms do not require "very small" θ to perform well. For instance, the competitive ratio of our mechanism for additive utility functions is $(1-1/e)\cdot(1-6\theta/5)$ when all the items have equal utilities [3,4] (Section E). To give examples, for $\theta = 1/20$ and $\theta = 1/40$ (which are reasonable assumptions in the applications that we will discuss) the competitive ratios are 0.592 and 0.613, respectively. We remark that if the Small Bidders assumption is dismissed, there exists no truthful mechanism with positive competitive ratio for the case of divisible items.[5]

Finally, we briefly explain whether having prior information about θ could help the buyer to design improved mechanisms. We construct a family of instances in which revealing θ to the buyer does not improve the competitive ratio of the optimal mechanism by more than $\theta/1.5$ (See Appendix A.1 for the details). This implies that, for typically small values of θ, knowing the value of θ is not in general significantly beneficial to the buyer. For relatively large values of θ, we conjecture that it is possible to design improved mechanisms when given some information about θ; this case is left open for future studies.

3.1. Roadmap

After a brief literature review in Section 4, we define the optimal mechanism in Section 5. Our mechanism is parameterized by a single input variable which we call a *standard allocation rule.* Most of the work in the later sections is about finding the optimal standard allocation rule and proving its optimality. In Section 6, we find the unique standard allocation rule which provides a mechanism with competitive ratio $1 - 1/e$. We complement this result in Section 7 by showing that no truthful mechanism attains a competitive ratio higher than $1 - 1/e$. Section 8 presents our results for indivisible items.

Our mechanisms for submodular utility functions are discussed in Section 9. We discuss the applications in Section 10. Section 11 is the conclusion.

4. Related Work

We study a simple and yet powerful framework that has been deployed in a wide range of applications, including allocating R&D subsidies by government agencies Ensthaler and Giebe (2014a,b), conducting auctions for reducing the emission of greenhouse gases Maskin (2002), Chung and Elly (2002), He and Chen (2014), and pricing tasks in crowdsourcing markets and other market places Singer (2010), Chen et al. (2011). We discuss these applications, their relation to our model, and some of the related work in our applications section, Section 10. Here, we highlight the main theoretical progress in the prior-free and Bayesian frameworks.

To the extent of our knowledge, the prior-free framework was first formalized by Singer (2010) under the name of *Budget Feasible Mechanism Design* framework. Singer (2010), followed by Chen et al. (2011), studied the prior-free model for additive and submodular utility functions. For the case of additive utilities and indivisible items, Singer (2010) designed a deterministic mechanism with competitive ratio $1/6$. Later, Chen et al. (2011) improved the competitive ratio to $1/(2+\sqrt{2})$, and also gave a randomized mechanism with competitive ratio $1/3$. For submodular utility functions, Singer (2010) gave a randomized mechanism with a competitive ratio of $1/112$, which was improved to $1/7.91$ by Chen et al. (2011).

The main idea followed by the previous work in the prior-free framework is designing *proportional share mechanisms* which use "cutoff allocation rules" and "linear payment rules". (e.g. see Singer (2010), Chen et al. (2011), Singer and Mittal (2013), Goel et al. (2014)) More precisely, these mechanisms typically allocate items for which the ratio of cost to utility is below a fixed threshold (computed by the mechanism), and pay to each seller an amount proportional to the utility of her item. We find that, under the Small Bidders assumption, the optimal allocation rule is in fact not a cutoff allocation rule, and also, its corresponding payment rule is non-linear (Section 5).

Maskin (2002) uses the Bayesian framework to analyze the UK emissions reduction auction. This work focuses on some special classes of distributions and finds optimal ex ante budget feasible mechanisms, i.e. the budget constraint has to be satisfied only in expectation. This result is generalized by Baron and Myerson (1982), who consider a larger class of mechanism design problems and translate them into the setting of Baron and Myerson (1982). Ensthaler and Giebe (2014a) characterize optimal ex ante budget feasible mechanisms and construct them for the class of regular cost distributions.

Jarman and Meisner (2015) study optimal ex post budget feasible mechanisms in a Bayesian setting. They show that the optimal mechanism in the class of all truthful *deterministic* mechanisms is a cut-off mechanism (i.e. could be represented by a set of cutoffs, one for each seller). They also construct such mechanisms when the cumulative cost distributions are log-concave. While noting that our model is prior-free and our benchmark for optimality is different, we remark that our optimal mechanism for indivisible items uses randomization.

5. The Optimal Mechanism

We present the optimal mechanism for divisible items in this section. To build ideas, we first introduce and formalize the notion of an *allocation rule*. An allocation rule $f : \mathbb{R}^+ \to [0, 1]$ is a function which determines how much to buy from a given seller. The domain of an allocation rule is the *cost-utility rate*; the allocation rule f declares that we should buy $f(\frac{c_i}{u_i})$ fraction of seller i's item. For notational convenience, we denote $\frac{c_i}{u_i}$ by d_i and call it the *cost-utility rate* of item i. We do not enforce using the same allocation rule for all sellers.

We say an allocation rule $f:\mathbb{R}_+\to[0,1]$ is a *Standard Allocation Rule* when f is continuous and strictly decreasing over $[0,e-1]$, $f(0)=1$, and $f(c)=0$ for $c\geq e-1$. (The choice of $e-1$ is merely for simplification of calculations; it can be replaced with any other constant.)

For any standard allocation rule f, we can define an associated family of allocation rules

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

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

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 pick the largest positive r such that $f_r\in\mathcal{F}(f)$ is "budget-feasible", meaning that 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 corresponding to allocation rule f_r; we define it next.

5.1. Payment Rule

Recall that given an allocation rule $f_r\in\mathcal{F}(f)$, the value of $f_r(d_i)$ only tells us what fraction of seller i's item to buy. But how much should we pay seller i to incentivize her to report her cost truthfully? We define these payments based on the well-known Myerson's characterization of truthful mechanisms Myerson (1981).

Given that we are using allocation rule f_r for seller i, let the payment rule for seller i be denoted by $P_{i,r}:\mathbb{R}_+\to\mathbb{R}_+$. In other words, let $P_{i,r}$ be the function that maps the reported cost of seller i to the payment it receives. To define $P_{i,r}$, we propose the following thought experiment: Divide seller i's item into u_i pieces of equal size (note that the seller's cost for each piece is d_i). Now, the allocation rule f_r can be seen as a function that maps the cost of a single piece into the fraction allocated from that piece. Let $Q_{f_r}(d):\mathbb{R}_+\to\mathbb{R}_+$ denote the function that maps the cost for a single piece into the payment for that piece. Now, by Myerson's characterization Myerson (1981), the payment for each piece is given by the following formula:

$$Q_{f_r}(d)=d\cdot f_r(d)+\int_d^{\infty} f_r(y)\,\mathrm{d}\,y.$$

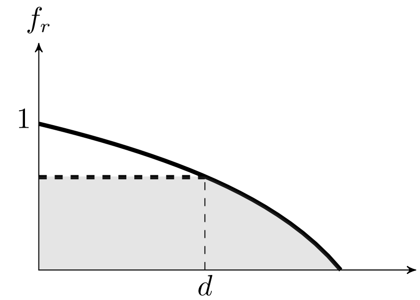

Figure 1 The shaded area under the curve represents $Q_r(d)$.

For national simplicity, from now on we denote Q_{f_r} by Q_r, when f_r is clearly known from the context. Intuitively, $Q_r(d)$ represents the area under the curve illustrated in Figure 1. Going forward, we will call the function Q_r a *unit-payment rule.*

Note that $P_{i,r}$ and Q_r are related by the equation

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

where c stands for the cost reported by seller i. Thus, to summarize, if we use an allocation rule f_r for seller i, then we buy $f_r(d_i)$ units of her item and pay her $P_{i,r}(c_i)$.

We can now formally define the notion of budget-feasibility of an allocation rule.

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

The notions of allocation rule, payment rule, and budget feasibility (defined above) are also explained in Section A.2 using an example.

5.2. The Optimal Mechanism

The basis of our optimal truthful mechanism is a non-truthful mechanism that we name Envy-Free(f) and is defined next. We call it Mechanism Envy-Free(f) since it satisfies an envy-freeness property: no seller who reports her true cost would envy another seller's payoff "per unit of utility". We defer the discussion on envy-freeness to Appendix A.3 since this property is not used in our other proofs.

Given any standard allocation rule f, Envy-Free(f) 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 budget-feasible; suppose this happens 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. We call r^* the *scaling ratio of the mechanism*. We define this process formally in Mechanism Envy-Free(f). Section A.2 contains an example that runs this mechanism on a simple instance.

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

---

input : Budget B, (u_i, c_i) 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^*;

Envy-Free(f) would have been a truthful mechanism if the allocation rule f_r that is offered to seller i had not depended on her private information (cost c_i in this case). 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 0 while keeping the cost of other sellers intact; then, we define r_i to be the scaling ratio of Mechanism Envy-Free(f) in this instance. The formal definition of the truthful mechanism appears in Mechanism Truthful(f).

We are now ready to present our main theorems. Let the standard allocation rule f^* be defined as $f^*(d) = ln(e-d)$ for $d \leq e-1$.

THEOREM 1. *Mechanism Truthful(f) is budget feasible and truthful. Moreover, it has competitive ratio $1-1/e$ when $f = f^*$*

THEOREM 2 **(Uniqueness)**. *The competitive ratio of Truthful(f) is strictly less than $1-1/e$ for any standard allocation rule $f \neq f^*$.*

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

```
input  : Budget B, (u_i, c_i) pair for each seller i
foreach i ∈ S do
    temp ← c_i;
    c_i ← 0;
    r_i ← Mechanism Envy-Free(f);
    c_i ← temp;
end
foreach i ∈ S do
    Allocate f_{r_i}(c_i/u_i) from seller i;
    Pay P_{i,r_i}(c_i) to seller i;
end
```

Theorem 3. *Any truthful budget feasible mechanism attains competitive ratio at most* $1 - 1/e$.

Theorems 1 and 3 together show that when $f = f^*$, Truthful(f) attains the highest competitive ratio within the class of truthful mechanisms. Theorem 2 shows that f^* is the unique allocation rule that attains this competitive ratio (For brevity, we call it the unique optimal allocation rule).

Theorems 1 and 2 are proved in Section 6. Theorem 3 is proved in Section 7. We can transform our mechanism for divisible items to an optimal mechanism for indivisible items; this is discussed in Section 8.

6. The (Unique) Optimal Allocation Rule

We present the proofs for Theorems 1 and 2. While Appendix C contains a self-contained independent proof for Theorem 1, this section presents a more intuitive proof under a simplifying assumption: we assume that $u_i = 1$, for all $i \in S$. Note that proving Theorem 2 under this assumption is without loss of generality. Furthermore, all the other results that we prove in this section (including the lemmas) still hold when this assumption is dismissed.

The role of this assumption is, essentially, simplifying the notation: when all the items have utility 1, the payment rule coincides with the unit-payment rule. This will make the analysis cleaner. Therefore, in this section we will use $f(c), Q_f(c)$ to denote an allocation rule and its corresponding payment rule, respectively; they determine allocation and payment for a seller who reports cost c.

We call an allocation rule f the *optimal allocation rule for* Truthful(f) if this mechanism attains the highest competitive ratio under f. Similarly, we say that f is the *optimal allocation rule for*

Envy-Free(f), if this mechanism attains the highest competitive ratio under f. (In computing the competitive ratio of Envy-Free(f), we always assume that sellers report their costs to it truthfully.) The analysis proceeds as follows. First, we show that in order to find the optimal allocation rule for Truthful(f), it suffices to find the optimal allocation rule for Envy-Free(f). Then, we find the optimal f for Envy-Free(f) in Section 6.1. The full proofs for all of the propositions, lemmas, and theorems in this section are presented in Appendix B, to avoid technical details. Here, we discuss the core components of the proofs.

Proposition 1. *The competitive ratio of Truthful(f) is no larger than the competitive ratio of Envy-Free(f). Moreover, these competitive ratios are equal when f is a log-concave function*[6].

We remark that the proof for the second part relies on the Small Bidders assumption. Intuitively, this assumption ensures that the market that remains after the removal of a seller (or after setting her bid to 0) is "sufficiently similar" to the original market. The Small Bidders assumption together with log-concavity of f will guarantee the scaling ratios r_i in Truthful(f) remain close to the scaling ratio r^* in Envy-Free(f), and so do the payments. This will imply that the competitive ratios are asymptotically equal. We will observe that f^* is log-concave, and this will be the only place that log-concavity affects our analysis.

Instead of proving Theorems 1 and 2, we can use Proposition 1 to prove their counterparts for Envy-Free(f). That is, to prove Theorems 1 and 2, it suffices to prove the following lemma.

Lemma 1. *The competitive ratio of Envy-Free(f) is $1 - 1/e$ for $f = f^*$, and is a constant strictly less than $1 - 1/e$ for any standard allocation rule $f \neq f^*$.*

Proof of Theorems 1, 2: By Lemma 1, the competitive ratio of Envy-Free(f) is $1 - 1/e$ for $f = f^*$ and is a constant strictly less than $1 - 1/e$ for any standard allocation rule $f \neq f^*$. Also, note that f^* is concave (and therefore, log-concave). Proposition 1 then implies that the competitive ratio of Truthful(f) is equal to the competitive ratio of Envy-Free(f) when $f = f^*$, and is at most equal to the competitive ratio of Envy-Free(f) otherwise. This proves both of the theorems. □

The rest of this section is devoted to proving Lemma 1. We take the first step by computing the competitive ratio of Envy-Free(f) for any f.[7] The core idea that derives the analysis is a mapping that maps any allocation rule f to another function $\widehat{f}$. We define $\widehat{f}$ such that

$$\widehat{f}(f(c)) = Q_f(c) - c, \ \forall c \in \text{supp}(f).$$

Note that $\widehat{f}$ is well-defined because f is a standard allocation rule, which means it is strictly decreasing, and, thus, invertible.[8] It is illuminating to represent $\widehat{f}$ in a two-dimensional plane that has the allocation rule f as the horizontal axis. Figure 2(a) illustrates this for a simple allocation rule $f(c) = 1 - c$ with $\text{supp}(f) = [0,1]$. This representation is useful because, interestingly, it exposes the competitive ratio of Envy-Free(f): For the case of $f(c) = 1 - c$, the competitive ratio would be equal to the x-intercept of $\widehat{f}$. (By x-intercept we mean distance of the origin from the point at which $\widehat{f}$ intersects the horizontal axis. See Figure 2(a).)

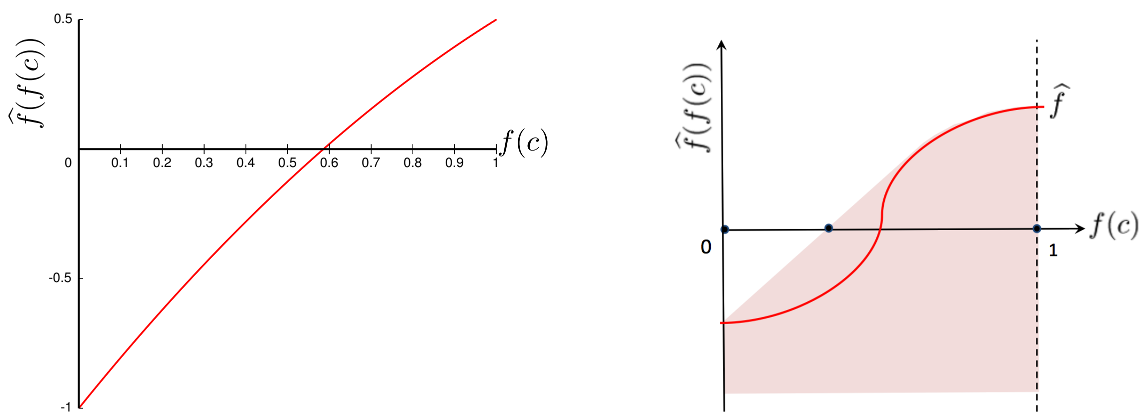

(a) The plotted curve represents $\widehat{f}(f(c))$ for $f(c) = 1 - c$. $\widehat{f}$ intersects the horizontal axis at ≈ 0.58.

(b) The x-intercept of $\text{Conv}^{\downarrow}\left\{\widehat{f}\right\}$ is marked in the picture for a hypothetical $\widehat{f}$.

Figure 2

In general, for any standard allocation rule f we can show that the competitive ratio of f is equal to the x-intercept of the lower convex hull of $\widehat{f}$. Figure 2(b) demonstrates this x-intercept graphically. More formally, define the epigraph and the lower convex hull of $\widehat{f}$ as

$$\text{epi}\left\{\widehat{f}\right\} = \left\{(f(c), y) : c \in \text{supp}(f), y \leq \widehat{f}(f(c))\right\},$$

$$\mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\} = \mathsf{Conv}\left\{\mathsf{epi}\left\{\widehat{f}\right\}\right\},$$

where $\mathsf{Conv}\{P\}$ denotes the convex hull of any set of points P. We prove that the competitive ratio of f is equal to the x-intercept of $\mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}$. It is worth pointing out that when $\widehat{f}$ is concave, the x-intercepts of $\widehat{f}$ and $\mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}$ coincide (which is the case in Figure 2(a)).

THEOREM 4. *For any standard allocation rule f, the competitive ratio of Envy-Free(f) is*

$$\inf\left\{x : (x,0) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}\right\}.$$

Using this theorem and Proposition 1 together, we can compute the competitive ratio of Truthful(f) for any log-concave f. (This is still possible even if f is not log-concave, but under a stronger version of the Small Bidders assumption that assumes $\max_{i\in S}\{\frac{c_i u_i}{B}\} \to 0$. See Appendix E, Propositions 5 and 6.)

Proof Sketch for Theorem 4: The proof for Theorem 4 has two parts: showing that the competitive ratio is *at least* what the theorem claims, and that it is *at most* what the theorem claims. Here, we sketch the proof for the first part.

Define $\alpha = \inf\left\{x : (x,0) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}\right\}$. By the separating hyper plane theorem, there should be a line l passing through $(\alpha, 0)$ such that all the points in $\mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}$ fall on the same side of l. Let $y = sx + b$ denote the line equation for l. We can show that $0 < s < \infty$. (See the complete proof.)

Now, suppose that Envy-Free(f) allocates $x_i = f(c_i)$ units from seller i. Then, we should have

$$\widehat{f}(x_i) \leq sx_i + b,$$

since $(x_i, \widehat{f}(x_i)) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}$. Summing up this inequality for all i implies that

$$\sum_{i=1}^{n} \widehat{f}(x_i) = \sum_{i=1}^{n} Q_f(c_i) - c_i \leq \sum_{i=1}^{n} sx_i + bn.$$

We are done if we have $\sum_{i=1}^{n} c_i \leq B$. In this case, the above inequality would imply that $0 \leq \sum_{i=1}^{n} sx_i + bn$, or equivalently, $\frac{-bn}{s} \leq \sum_{i=1}^{n} x_i$. Now, note that on the left-hand side, $n = U^*$ and $\frac{-b}{s}$ is exactly the x-intercept of l. In the other hand, the right-hand side is just the total utility attained by the mechanism. So, the theorem is proved when $\sum_{i=1}^{n} c_i \le B$ holds. The case $\sum_{i=1}^{n} c_i > B$ needs an additional trick that involves writing stronger versions of the above inequalities; we address it in the complete version of the proof in the appendix. □

6.1. Proof Sketch for Lemma 1.

We are now ready to prove Lemma 1. In this section we use $\widehat{f}$ to denote the set

$$\widehat{f} = \{(f(c), Q_f(c) - c) : \forall c \in \mathrm{supp}(f)\}.$$

We still use $\widehat{f}(x)$ to denote the unique y for which $(x, y) \in \widehat{f}$.

The next observation plays an important role in the proof.

Proposition 2. *$\widehat{f^*}$ corresponds to a straight line passing through the point $(1-1/e, 0)$ with positive slope. Moreover, any line that passes through $(1-1/e, 0)$ with positive slope corresponds to $\widehat{f_r^*}$ for some $r > 0$, and vice versa.*

Figure 3 is an informal proof-by-picture for this proposition. In Section H in the appendix, we provide more intuition on why $\widehat{f^*}$ should indeed be a straight line. Our discussion there also offers more intuition on the definition of $\widehat{f}$.

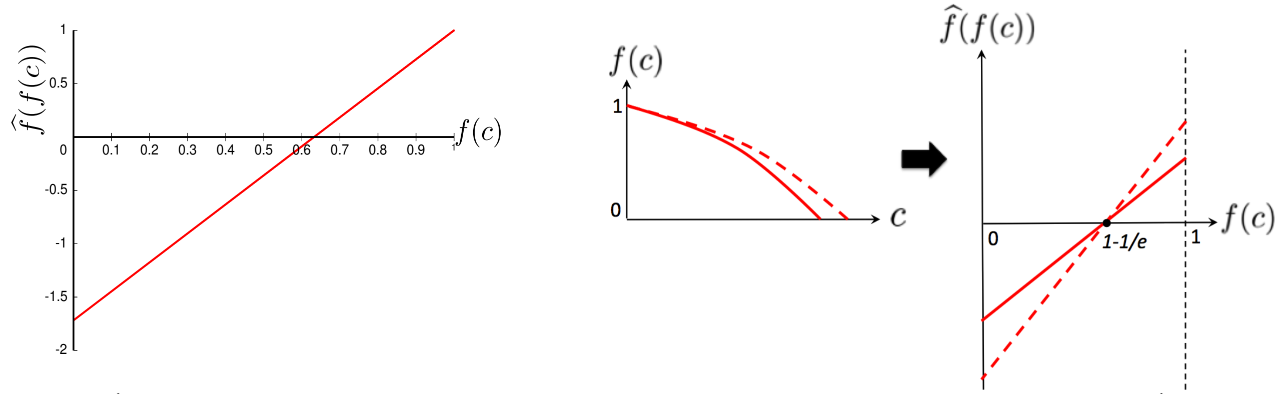

(a) $\widehat{f}$ is a straight line when $f = f^*$.

(b) Scaling of f horizontally corresponds to scaling $\widehat{f}$ vertically.

Figure 3

Proposition 2 says that $\widehat{f^*}$ is a straight line with x-intercept $1-1/e$. Together with Theorem 4, they imply that Mechanism Envy-Free(f^*) has competitive ratio $1-1/e$. This proves a part of

Lemma 1. To complete the proof of Lemma 1, we need to show that the competitive ratio of Envy-Free(f) is strictly less than $1-1/e$ for any standard allocation rule $f \neq f^*$. We prove this in Lemma 3, after stating a useful lemma.

LEMMA 2. *For any two allocation rules f, g, $\widehat{f} \cap \widehat{g} \neq \emptyset$.*

LEMMA 3. *For any standard allocation rule $f \neq f^*$,*

$$\inf\left\{x : (x,0) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}\right\} < 1-1/e.$$

Proof Sketch for Lemma 3: We prove a weaker version of this lemma in which the strict inequality in the lemma statement is replaced with a non-strict inequality. This weak version implies that f^* is an optimal allocation rule for Envy-Free(f), but it does not show its uniqueness. The proof for uniqueness is in the same spirit, but is more subtle; it is relegated to the appendix.

Suppose there exists an allocation rule g for which Mechanism Envy-Free(g) has a competitive ratio larger than $1-1/e$. By Theorem 4, we should have

$$1-1/e < \inf\{x : (x,0) \in \mathsf{Conv}^{\downarrow}(\widehat{g})\}.$$

Therefore, there should be a straight line with positive slope that passes through $(1-1/e, 0)$ and does not intersect $\widehat{g}$. By Proposition 2, this line corresponds to $\widehat{f_r^*}$, for some $r > 0$ (see Figure 4). Consequently, $\widehat{f_r^*} \cap \widehat{g} = \emptyset$, which contradicts Lemma 2. □

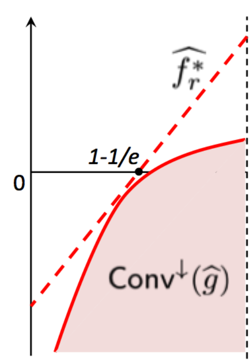

Figure 4 Choosing a positive r such that $\widehat{f_r^*}$ does not intersect $\mathsf{Conv}^{\downarrow}\{\widehat{g}\}$.

7. Impossibility Result

In this section we show that no truthful (and possibly) randomized mechanism achieves competitive ratio higher than $1-1/e$. We prove a stronger claim by allowing the mechanism to be budget-feasible 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 competitive ratio better than $1-1/e$; this also implies that no mechanism for the prior-free setting can achieve ratio better than $1-1/e$. [9] The proof is done in two steps. First, we show that given any truthful mechanism for this instance, there exists a simple posted price mechanism that achieves at least the same utility (in expectation). 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 no choice of p this mechanism can achieve a ratio better than $1-1/e$. The proof that we present, without loss of generality, analyzes the market in expectation: budget feasibility is satisfied in expectation; also, the utility derived by a mechanism is computed in expectation.

We now proceed to the full proof by first defining our impossibility 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 distribution function F, defined as follows:

$$F(c) = \frac{1}{e(1-c)}, \quad \forall c : 0 \leq c \leq 1-1/e. \tag{2}$$

In other words, $F(c)$ denotes the probability that the cost of a seller is at most c. Let $\mathcal{D}$ be the distribution defined by F and let $\overline{c}$ denote the expected cost of a seller sampled from $\mathcal{D}$, i.e. $\overline{c} = \mathbb{E}_{c\sim\mathcal{D}}[c]$. Define the budget to be $B = \overline{c} \cdot n$ where n denotes the number of sellers. We do the analysis assuming that n approaches infinity.

Definition 3. 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 4. A *uniform posted price mechanism* is a posted price mechanism that offers the same price to all sellers.

Definition 5. A cutoff allocation rule is an allocation rule which allocates the whole unit of an item if its cost is less than a certain cutoff and allocates nothing (0 units) of that item otherwise. Let $\mathsf{cutoff}(p)$ denote a cutoff allocation rule with the cutoff price p.

By definition, posted price mechanisms use cutoff allocation rules for allocating items. Note that we distinguish between the notions of allocation rule and selection rule. Recall from Section 2 that selection rule $A_i : (\mathbb{R}_+)^n \to [0,1]$ is a function that takes as its input the costs of *all* sellers and outputs the fraction of seller i's item that buyer buys. An allocation rule is different in that it takes the cost of a single seller as its input. Therefore, $A_i(\cdot, \mathbf{c}_{-i}) : \mathbb{R}_+ \to [0,1]$ is an allocation rule.

Lemma 4. *If the sellers costs are drawn i.i.d. from the distribution $\mathcal{D}$, then for any truthful mechanism there exists a posted price mechanism with the same competitive ratio.*

Proof. Due to Myerson's Lemma (Myerson (1981)), any truthful mechanism in Bayesian setting can be represented as a selection rule A paired with a payment rule P (recall the definitions of A, P from Section 2). The selection rule for seller i, A_i, is decreasing in c_i, and the payment rule for seller i, P_i, is defined according to the allocation rule $A_i(\cdot, \mathbf{c}_{-i})$ as we saw in Figure 1. A simple way to implement the allocation rule $A_i(\cdot, \mathbf{c}_{-i})$ is by finding a distribution π_i over cutoff allocation rules.[10] In other words, the distribution π_i is a distribution over cutoff prices, and can be interpreted in the following way: first sample a price p from π_i, and then offer that price to seller i, i.e. use the allocation rule $\mathsf{cutoff}(p)$ to buy from i.

We prove the existence of π_i by constructing it.

Claim 1. *Define the distribution π_i by its CDF $F_i(\cdot)$ such that $F_i(c) = 1 - A_i(c, \mathbf{c}_{-i})$ for all costs $c \geq 0$. Then, π_i would implement $A_i(\cdot, \mathbf{c}_{-i})$.*

Proof. All we need to show is that, when a cutoff price is sampled from π_i, then the item of a seller with cost c is bought with probability $A_i(c, \mathbf{c}_{-i})$. To see this, note that the probability that the item is bought is exactly equal to the probability of sampling a cutoff price at least c. This probability is equal to $1-(1-A_i(c,\mathbf{c}_{-i}))$ by the definition of π_i; this proves the claim. □

Let $f_i(\cdot)$ denote the PDF of π_i. We claim that the cutoff allocation rule with the cutoff price

$$p_i = F^{-1}\left(\int_0^\infty f_i(p)\cdot F(p)dp\right) \tag{3}$$

in expectation attains the same utility and spends the same budget as the allocation rule $A_i(\cdot,\mathbf{c}_{-i})$ paired with its corresponding Myerson payment rule. (Recall that F is defined by (2).)

CLAIM 2. *For any seller $i \in S$, cutoff(p_i) achieves the same utility and spends the same budget (in expectation) as the allocation rule $A_i(\cdot,\mathbf{c}_{-i})$ paired with its corresponding Myerson payment rule.*

Proof. The main idea of the proof is that the set of points $P = \{(F(p), pF(p)) : 0 \le p \le 1-1/e\}$ forms a straight line segment in the two-dimensional plane; see Figure 5 for a proof by picture. Note that $F(p)$ denotes the expected allocation when a price p is offered to a seller and $pF(p)$ is

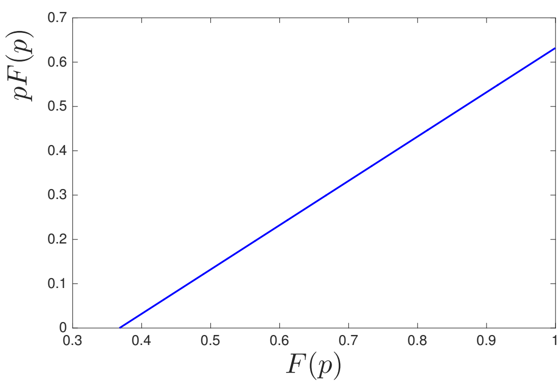

Figure 5 This parametric plot is representing the set of points P for $p \in [0, 1-1/e]$.

the corresponding expected payment.

Now, see that the expected utility achieved by the allocation rule $A_i(\cdot,\mathbf{c}_{-i})$ is $\int_0^\infty f_i(p)F(p)dp$, which is exactly equal to the expected utility achieved by cutoff(p) due to (3). To prove the claim, it remains to verify that the Myerson payment rules corresponding to $\mathsf{cutoff}(p_i)$ and $A_i(\cdot, \mathbf{c}_{-i})$ spend the same budget (in expectation). To this end, observe that P is a straight line; consequently, since the allocation rules $A_i(\cdot, \mathbf{c}_{-i})$ and $\mathsf{cutoff}(p_i)$ have the same expected allocation, they also spend the same amount of budget in expectation. □

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 lemma. □

LEMMA 5. *If the sellers' costs are drawn i.i.d. from the distribution $\mathcal{D}$, then for any posted price mechanism there exists a uniform posted price mechanism with the same competitive ratio and the same expected payment.*

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

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

First, observe that the uniform posted price mechanism with price $\overline{p}$ achieves a utility equal to the utility of the original posted price mechanism; this is simply because $F(\overline{p})$, the expected utility gained by offering the price $\overline{p}$ to a seller, is the average of $F(p_i)$ where the average is taken over all $i \in S$. It remains to verify that the uniform posted price mechanism has the same expected payment. To this end, just observe that the set P is a straight line (depicted in Figure 5). Consequently, since the posted price mechanism and the uniform posted price mechanism have the same expected allocation, they also spend the same amount of budget in expectation. □

Next is the main theorem of this section. We prove that any truthful mechanism for indivisible items has competitive ratio at most $1 - 1/e$. We will see that the counterpart for divisible items would be implied as a corollary.

THEOREM 5. *Any truthful mechanism for indivisible items that is budget-feasible in expectation has competitive ratio at most $1 - 1/e$.*

Proof. We use Lemma 5 and show that no uniform posted price mechanism can attain a competitive ratio better than $1-1/e$. Equivalently, we show that the optimum uniform posted price mechanism, i.e. the mechanism which spends all the budget in expectation, has competitive ratio at most $1-1/e$.

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 our definitions for $F(\cdot)$ and B, we can solve this equation to get $p^* = \frac{e-2}{e-1}$. Now, we are ready to compute the competitive ratio. First, note that the (expected) utility of the uniform posted price mechanism is $n\cdot F(p^*)$. Given this fact, the theorem could be easily proved if we had $\sum_{i\in S} c_i \le B$: Then, we would have had $U^* = n$ (the optimum solution could buy all items), and so we could write the competitive ratio as

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

which would prove the claim. However, although budget feasibility is satisfied in expectation, i.e. $\mathbb{E}\left[\sum_{i\in S} c_i\right] = B$, the sum is not always bounded by B, which means $U^* = n$ does not always hold. We can fix this issue using Hoeffding bounds (see Section J 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 J), for any $\epsilon > 0$ we have:

$$\Pr\left[\sum_{i\in S} c_i \ge (1+\epsilon)\cdot B\right] \le e^{-\Omega(|S|)} \tag{4}$$

Recall that we defined $n = |S|$ and that in our hardness instance $n\to\infty$. Using (4), we will provide an upper bound on the competitive ratio which, for any constant $\epsilon > 0$, approaches to $(1-1/e)(1+\epsilon)$ as n approaches infinity. This would prove that the competitive ratio is a constant not larger than $1-1/e$.

To this end, first note that if $\sum_{i\in S} c_i \le B(1+\epsilon)$, then we have $U^* \ge \frac{n}{1+\epsilon} - 1$; this holds due to Lemma 6 (the -1 in the right-hand side is required due to indivisibility of the items). We can use this fact along with (4) to write the following upper bound on the (expected) competitive 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 $n \to \infty$. Recall that $F(p^*) = 1 - 1/e$; this finishes the proof. □

Now we use Theorem 5 to prove its counterpart for divisible items.

COROLLARY 1. *Any truthful mechanism for divisible items that is budget-feasible in expectation has competitive ratio at most* $1 - 1/e$.

Proof. Proof by contradiction. Suppose there exists such a mechanism with competitive ratio $\alpha > 1 - 1/e$ for some constant α. Then, we show that we can convert this mechanism to an α-competitive mechanism for indivisible items which is truthful and budget-feasible in expectation. This would contradict Theorem 5.

To do this conversion, we repeat the exact same argument that we used to prove Theorem 5. Using the same argument, we can convert the given α-competitive mechanism to a uniform posted price mechanism with competitive ratio α. Note that all posted price mechanisms allocate items without dividing them. Consequently, we have an α-competitive mechanism for indivisible items. Contradiction. □

We discuss our impossibility instance further in Section H in the appendix. We rediscover the distribution given by (2) in a stylized Bayesian model, and offer an alternative explanation of why this distribution should be mapped to a straight line in Figure 5. This discussion could be interesting to readers familiar with the Myersonian approach in revenue maximizing forward auctions.

8. Mechanisms for Indivisible Items

In this section, we explain how our mechanism for divisible items can be converted to a mechanism for indivisible items; the resulting mechanism would have the same competitive ratio, $1 - 1/e$.

When items are indivisible, a natural approach is to find a single cutoff r that determines the allocation and payment: the item of seller i is allocated iff $\frac{u_i}{c_i} \leq r$, and the seller will be paid ru_i. The cutoff r is selected to be the largest number for which the payments sum to be at most B. This simple mechanism captures the main idea behind the "proportional share" mechanisms used in most of the previous studies, such as Singer (2010), Chen et al. (2011), and Ensthaler and Giebe (2014b). (They do not consider the Small Bidders assumption.)

We can show that this mechanism is individually rational, truthful, and budget-feasible, and that it has competitive ratio $\frac{1}{2}$ under the Small Bidders assumption.[11] We remark that this simple mechanism is in fact identical to Mechanism Envy-Free($\bar{f}$) where $\bar{f}$ is a *step function*, i.e. $\bar{f}(d)=1$ for $d \leq 1$, and $\bar{f}(d)=0$ for $d>1$. This is the only allocation rule known to us for which Envy-Free(f) could be truthful. With a proper tiebreaking rule, this mechanism is truthful, budget feasible, and has competitive ratio 1/2 under the Small Bidders assumption (as shown in Appendix H.1). This is also demonstrated in Figure 6: recall Theorem 4 which characterizes the competitive ratio of Envy-Free(f) for any standard allocation rule f as the x-intercept of the lower convex hull of $\widehat{f}$. Although $\bar{f}$ is not a standard allocation rule (since it is not strictly decreasing), it is possible to construct a sequence of standard allocation rules that converges to f. Figure 6 demonstrates that the corresponding x-intercepts in such a sequence converge to 1/2.

So, not only does the mechanism always attain at least 1/2 of the optimum utility, but also there are instances for which the mechanism attains exactly 1/2 of the optimum utility. We will discuss such an instance at the end of this section. (It is also possible to construct such an instance directly from Figure 6; see Appendix B, proof of Lemma 11 for this alternative approach.)

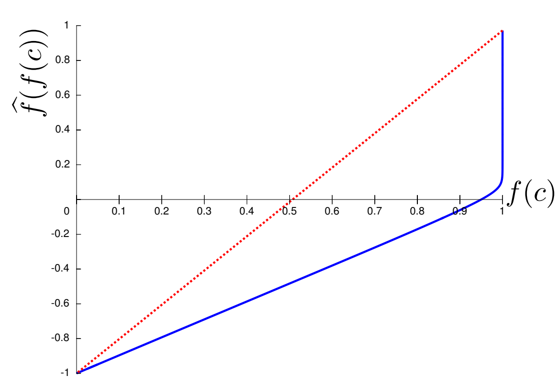

Figure 6 The solid curve represents $\widehat{f}(f(c))$ for $f \approx \bar{f}$. (Approximation uses the logistic function.) The dotted line represents the border of the lower convex hull, which has x-intercept ≈ 0.5.

We can design a mechanism with a higher competitive ratio using randomization. The idea is to first treat items as divisible and run the mechanism for divisible items. Then, we convert the obtained (fractional) allocation to an integral allocation for indivisible items. We design a lottery (rounding process) that takes the fractional allocation as its input and outputs an integral allocation with its corresponding payments. By the design of our lottery, the resulting mechanism is individually rational, truthful, and budget-feasible. Also, its competitive ratio is (asymptotically) equal to $1 - 1/e$.

THEOREM 6. *Mechanism Truthful(f^*) can be converted to a mechanism for indivisible items which remains individually rational, truthful and budget feasible and has competitive ratio* $1 - 1/e$.

This theorem is proved in Appendix G. The proof uses the Small Bidders assumption. Under this assumption, we are able to guarantee that the solution for indivisible items attains asymptotically the same level of utility as the solution for divisible items, while guaranteeing ex post budget feasibility. By Theorem 5, this is the highest possible competitive ratio for indivisible items. Remarkably, it is still unknown to us whether deterministic mechanisms could attain the same competitive ratio for indivisible items.

We end this section with an example which provides some intuition on how randomization helps to improve the competitive ratio of the proportional share mechanism. Our example contains $2n$ sellers, n of which have items with cost 0 (low-cost sellers) and the other n have items with cost 1 (high-cost sellers). All items provide utility 1 to the buyer, and the buyer's budget is n. Let $n \to \infty$; this ensures that the Small Bidders assumption holds.

Recall that in the proportional share mechanism, a single cutoff r determines the payment to each seller: the payment is r per unit of utility to winners and 0 to losers. In our example, the mechanism allocates the items of sellers with cost 0 and pays 1 to each, having no budget left to buy from high-cost sellers. Could we somehow pay less to low-cost sellers? With randomization, the answer is positive. Randomization allows us to buy some "probability shares" at a rate lower than r from low-cost sellers, at the expense of buying some probability shares from high-cost sellers at a rate higher than r. In other words, randomization provides the option of paying less to low-cost sellers in exchange for paying more to high-cost sellers (per unit of utility). Exploiting this tradeoff improves the mechanism.[12] In Sections H.2 and H.3, we study this tradeoff further in a stylized Bayesian setting and rediscover the allocation rule f^*. This discussion could be insightful to readers familiar with the Myersonian approach in revenue maximizing forward auctions.

Figure 7 compares the outcomes of Mechanisms Envy-Free($\bar{f}$) and Envy-Free(f^*) in our example. (We focus on the outcome of Envy-Free(f^*), and not Truthful(f^*), for its simplicity; note that these mechanisms have the same competitive ratio by Proposition 1.) The scaling ratios in these mechanisms are respectively $r^* \approx 0.70$ and $\overline{r} = 1$. The vertical axis represents the probability shares bought from sellers. Envy-Free($\bar{f}$) buys the probability shares only from low-cost sellers and pays them at rate 1 (dollar per probability share). Envy-Free(f^*) buys probability shares from both groups. After Envy-Free(f^*) buys the first $f^*_{r^*}(1)$ units of probability shares from low-cost sellers, it pays them at a rate lower than 1 for the rest. Because of this, the mechanism spends the budget more efficiently overall and allocates more utility than Envy-Free($\bar{f}$).

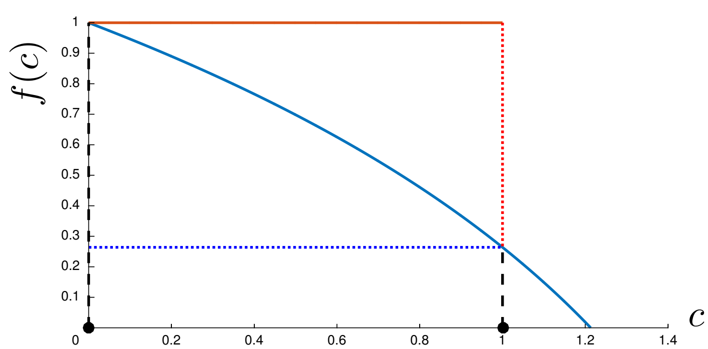

Figure 7 The allocation rules $f^*_{r^*}$ (blue curve) and $\bar{f}_{\overline{r}}$ (red curve). Mechanisms Envy-Free(f^*) and Envy-Free($\bar{f}$) allocate 1 unit from each low-cost seller, and respectively $f^*_{r^*}(1)$ and 0 units from each high-cost seller.

Next, we discuss how to transform the fractional allocation generated by Envy-Free(f^*) to an integral allocation. We give a simple sketch; the general argument is in Appendix G. Suppose that in the fractional solution, $x_L (= 1)$ and x_H respectively denote the fraction allocated from low-cost and high-cost sellers, and let p_L, p_H be the corresponding payments. Construct the integral solution as follows: allocate 1 unit from each low-cost seller and pay her p_L. Also, choose $x_H \cdot n$ of the high-cost sellers uniformly at random, allocate 1 unit from each, and pay each p_H/x_H . (We assume that $x_H \cdot n$ is an integer for the sake of simplicity.) This mechanism satisfies the desired properties: its payments sum up to at most B, and its competitive ratio is equal to the competitive ratio of Envy-Free(f^*).[13] The mechanism remains truthful and ex post individually rational.

9. Submodular Utility Functions

This section contains a summary of our results for when the utility function of the buyer is a *monotone* submodular function rather than an additive function; Section I contains the complete discussion.

Suppose the utility function of the buyer is represented by a monotone submodular function $F : 2^S \to \mathbb{R}_+$, i.e. $F(T)$ represents the utility of a subset $T \subseteq S$ for the buyer. The buyer's problem then becomes selecting a budget feasible subset of sellers with maximum possible utility. In other words, the optimum subset for the buyer is a subset S^* which is budget feasible ($\sum_{i \in S^*} c_i \leq B$) and has the highest utility among all the budget feasible subsets.

This problem was first studied in Singer (2010) without the Small Bidders assumption and a 0.0089-competitive mechanism was presented for it. Later, Chen et al. (2011) improved this result by giving an exponential-time deterministic mechanism with competitive ratio 0.119 and a polynomial-time randomized mechanism with competitive ratio 0.126. We study this problem under the *Small Bidders assumption* (see Section 2.3 or I.1 for formal definitions) and provide more efficient mechanisms for this case. A summary of our results appears below.

Our Contributions

All of our results hold under the Small Bidders assumption.

1. In Section I.2, we design a deterministic mechanism with competitive ratio $\frac{1}{2}$, which may have an exponential running time.

2. We show that the above mechanism has a polynomial running time and competitive ratio $\gamma^2/2$ when it has access to a γ-approximation oracle for solving the corresponding knapsack optimization problem (see Section I.2 for details). To the extent of our knowledge, the best existing oracle has approximation ratio $\gamma = 1 - 1/e$ due to Sviridenko (2004); this provides us a polynomial-time mechanism with competitive ratio $\gamma^2/2 \approx 0.2$.
3. In Section I.3, we improve the above result by presenting a deterministic polynomial-time mechanism with competitive ratio $\frac{1}{3}$. The mechanism that we use here is similar to the the mechanism used in Singer (2010), adapted to work under the Small Bidders assumption.

It is worth pointing out that the exponential running time of our $\frac{1}{2}$-competitive mechanism is solely due to the computational complexity of solving the corresponding knapsack optimization problem. The details are fully discussed in Section I. We remark that our $1 - 1/e$ upper bound on the competitive ratio of any truthful mechanism (Theorem 3) still holds here. Interestingly, this upper bound matches the computational complexity upper bound for the underlying optimization problem for submodular functions (which is $1 - 1/e$, by Sviridenko (2004)), although they are induced by seemingly very different arguments.

10. Applications

The main applications that we discuss in this section are allocation of R&D subsidies by government agencies, conducting auctions for reducing the emission of greenhouse gases, and pricing tasks in crowdsourcing markets. In the end, we briefly mention how our work improves the pricing mechanisms designed for two other market places. Through out the discussion of each application, we also highlight some of the related work.

10.1. Allocation of R&D Subsidies by Government Agencies.

Subsidization of private R&D by granting public funds has received considerable attention in the literature (see David et al. (1999) for an extensive discussion). In addition to the benefits gained from financing socially valuable projects and encouraging innovation by small firms or startups, the benefits from signaling effect of subsidies has also been explained theoretically Takalo and Tanayama (2008) and studied empirically Meuleman and De Maeseneire (2012). Because of such benefits[14], direct subsidization of private R&D has received significant attention. For instance, "The small business Innovation Research (SBIR) program in the United States provides funds in excess of $1 billion annually to encourage innovation by small and medium-sized private enterprises" Ensthaler and Giebe (2014b). Another example is funding researchers who apply for grants by submitting a detailed plan of research and an associated cost. After evaluation of the proposals, offers are made to a subset of researchers. In both of these examples, the procurer's objective is maximizing total welfare (which we model as the sum of utilities) subject to a budget constraint.

The problem of subsidizing R&D activities naturally nests in our model: The Small Bidders assumption is reasonable in this setting, e.g. recall the $1 billion annual funds of SBIR program for small and medium-sized private enterprises Ensthaler and Giebe (2014b). Moreover, the prior-free setting is likely to be more practical than the Bayesian setting: The costs highly depend on the business specifics, and although sellers (i.e. businesses in this example) could have reasonable estimates of their own costs, it could be unreasonable to assume that the buyer knows the cost distribution for each seller. In addition, bidders do not bid repeatedly, which means sensible empirical distributions are hardly available.

When the cost distributions are known, the results developed for the Bayesian setting are applicable. Under some mild assumptions on cost distributions, Ensthaler and Giebe (2014a) and Jarman and Meisner (2015) find ex ante and ex post budget feasible mechanisms, respectively.

10.2. Emission Reduction Auctions

Subsidizing emission-reducing activities is used by both developed and developing countries to control greenhouse gas emissions. Conducting auctions is a practiced method for allocation of such subsidies (e.g. see Maskin (2002), Chung and Elly (2002)). In such auctions, the government spends a predetermined fixed budget and pays firms to limit their emissions, with the objective of maximizing the emission reduction. The bid of each firm in the auction contains a per-unit-reduction cost and a maximal reduction capacity. The government's objective is maximizing the emission reduction, subject to its budget constraint. This scenario fits our model for divisible items. The prior-free framework could be more practical here because firms' bids depend on various evolving factors such as their market power, their competitors' market power, and their technology, which could be unknown to the mechanism designer. Therefore, finding a reasonable empirical distribution is difficult. The prior-free approach is robust to such uncertainties. Also, our Small Bidders assumption holds when the government's budget is significantly larger than the total cost of each firm, which is a reasonable assumption for auctions of this scale.[15]

In the other hand, if the distributions of the firms' bids are known, and when these distributions satisfy the required assumptions, then the results provided for the Bayesian framework are applicable. Maskin (2002), followed by Chung and Elly (2002), analyzes the UK emission reduction auction and derives optimal ex ante budget feasible mechanisms for special classes of bid distributions. He and Chen (2014) analyze a similar auction proposed by Maskin (2011). The close connection between Maskin's auctions and the budget feasible mechanism design framework is also noted in Ensthaler and Giebe (2014a).

10.3. Pricing Tasks in Crowdsourcing Markets.

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 such as image annotation, data labeling for machine learning systems, consumer surveys, rating search engine results, spam detection, and product reviews. 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 pricing the tasks properly. Pricing the tasks too low can disincentivize workers to work on the tasks, while pricing too high results in a less efficient outcome for the requester (the one who procures workers). A natural approach to prevent economic loss from poor pricing is designing direct revelation mechanisms that solicit bids from workers and decide which workers to procure and how much to pay them. These mechanisms typically should ensure truthfulness, budget feasibility, and individual rationality. The main idea followed by the previous work (e.g. Singer and Mittal (2013), Singla and Krause (2013), Goel et al. (2014)) is designing proportional share mechanisms. Such mechanisms are sub-optimal if the Small Bidders assumption holds (see Section 8).

Consider a buyer with a limited budget and a large crowdsourcing task that requires multiple workers. One of the market places for such tasks is, e.g., CrowdFlower (CrdFlwr). The budget feasible mechanism design framework could be used to procure workers. For example, when workers are paid in an hourly basis, our results for divisible items are applicable: the items model workers' availability time. Our Small Bidders assumption holds when the ratio of each worker's cost to the total budget is small. This holds, for instance, when workers in the consideration set do not have very different reservation prices. A typical example of tasks meeting this criteria is annotation of (a subset of) the images in a large database with hourly payments.

10.4. Applications with non-additive utilities

Maximizing influence in social networks. Singer (2011), motivated by marketing applications, studies the problem of identifying a small subset of individuals in a social network that can serve as early adopters of a new technology and trigger a large cascade in the network. Given a limited budget, the goal is to maximize the "influence" (the word-of-mouth effect) by paying a subset of agents to be the initial adopters of the technology. Assuming that each node in the network has a private reservation price for being an initial adopter, Singer (2011) takes a mechanism design approach. Under the budget feasible mechanism design framework, he provides a truthful mechanism with competitive ratio ≈ 0.032. Below we explain how our results in Section 9 provide mechanisms with competitive ratios $\frac{1}{2}, \frac{1}{3}$ (for exponential and polynomial running time, respectively).

In the influence maximization model studied by Singer (2011), the “influence function” is proven to be a submodular function Kempe et al. (2003). Also, our Alternative Small Bidders assumption holds when each individual has a small effect on the total influence, which is sensible in social networks.[16] With this assumption, our results for submodular utility functions are applicable.

Experiment design. Horel et al. (2013) study a classic experiment design problem through the budget feasible mechanism design framework. The experimenter is given a fixed budget and has access to population of potential experiment subjects with private reservation prices. The experimenter’s goal is conducting a budget feasible subset of experiments that maximizes the “information gain”. Horel et al. (2013) show that the objective function corresponding to information gain is submodular. Then, they use this fact to design an approximately truthful budget feasible mechanism with competitive ratio ≈ 0.077. Our results in Section 9 provide truthful mechanisms with competitive ratios $\frac{1}{2}, \frac{1}{3}$ for this application. Note that our Alternative Small Bidders assumption fits well: it says that each subject has a small effect on the information gain.

11. Conclusion

The budget feasible mechanism design literature has often used cutoff allocation rules and their corresponding linear payment rules. We revisit the budget-feasible mechanism design framework under the Small Bidders assumption. Our main contribution is designing optimal budget feasible mechanisms under the Small Bidders assumption for when the utility function is additive. Interestingly, we find that the optimal mechanism uses a non-linear payment rule corresponding to a non-cutoff allocation rule.

We also find the optimal mechanism for the case of indivisible items. The idea is running the mechanism for divisible items, and then *implementing* the obtained fractional allocation. We design an implementation mechanism that takes the fractional allocation as its input and outputs an integral allocation with the corresponding payments. The resulting mechanism is individually rational, truthful, and budget-feasible; also, it attains the optimal competitive ratio, $1 - 1/e$. We provide a matching lower bound by showing that no budget feasible and truthful mechanism, for divisible or indivisible items, can attain a larger competitive ratio. In fact, we prove a stronger statement by showing that the lower bound holds for all truthful mechanisms which are budget feasible in expectation.

Finally, we study the problem for submodular utility functions. We first design a deterministic mechanism with competitive ratio $\frac{1}{2}$; this mechanism can have an exponential running time in general. Inspired by this mechanism, we also design a polynomial-time deterministic mechanism with competitive ratio $\frac{1}{3}$. These results only hold for indivisible items. To provide counterparts for divisible items (which could also lead to better mechanisms for indivisible items) one has to extend the domain of the submodular utility function to all possible allocations (the n-dimensional hypercube). The multilinear extension (Vondrak (2008)) or Lovàsz extension of submodular functions are potential choices for this purpose. We leave this case open for future study.

References

Amazon's mechanical turk platform. https://www.mturk.com/.

Crowdflower. http://www.crowdflower.com/.

Baron, D. P. and Myerson, R. B. 1982. Regulating a monopolist with unknown costs. *Econometrica: Journal of the Econometric Society*, 911–930.

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.

Chung, K.-S. and Elly, J. C. 2002. Mechanism design with budget constraint. Mimeo.

David, P. A., Hall, B. H., and Toole, A. A. 1999. Is public R&D a complement or substitute for private R&D? a review of the econometric evidence. Working Paper 7373, National Bureau of Economic Research. October.

Ensthaler, L. and Giebe, T. 2014a. Bayesian optimal knapsack procurement. *European Journal of Operational Research 234,* 3, 774 – 779.

Ensthaler, L. and Giebe, T. 2014b. A dynamic auction for multi-object procurement under a hard budget constraint. *Research Policy 43,* 1, 179 – 189.

Goel, G., Nikzad, A., and Singla, A. 2014. Matching workers expertise with tasks: Incentives in heterogeneous crowdsourcing markets. *Conference on Human Computation & Crowdsourcing (HCOMP)*.

GreenInitiative. 2015. Regional greenhouse gas initiative. https://www.rggi.org/.

He, H. and Chen, Y. April 2014. Auction mechanisms for allocating subsidies for carbon emissions reduction: An experimental investigation. *Resources For the Future discussion paper EfD 14-06*.

Horel, T., Ioannidis, S., and Muthukrishnan, S. 2013. Budget feasible mechanisms for experimental design. *CoRR abs/1302.5724*.

Jarman, F. and Meisner, V. 2015. Ex-post optimal knapsack procurement. Working Papers 15-02, University of Mannheim, Department of Economics.

Kempe, D., Kleinberg, J., and Tardos, E. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '03. ACM, New York, NY, USA, 137–146.

Maskin, E. 2002. How to reduce greenhouse gas emissions-an application of auction theory. *Nancy L. Schwartz Memorial Lecture, Northwestern University 41,* 3, 580–591.

Maskin, E. 2011. Notes on auctions for pollution reduction. *Keynote Speech at the 18th annual conference of European Association for Environmental and Resource Economists, Rome*.

Mehta, A., Saberi, A., Vazirani, U. V., and Vazirani, V. V. 2007. Adwords and generalized online matching. *J. ACM 54,* 5.

Meuleman, M. and De Maeseneire, W. 2012. Do R&D subsidies affect smes' access to external financing? *Research Policy 41,* 3, 580–591.

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. 2013. 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.

TAKALO, T. AND TANAYAMA, T. 2008. Adverse selection and financing of innovation: is there a need for R&D subsidies? Research Discussion Papers 19/2008, Bank of Finland. Sept.

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.

Appendix A: Miscellaneous Discussions from Sections 3 and 5

A.1. Revelation of θ to the Buyer

In Section 3, we mentioned that revelation of θ to the buyer cannot improve the competitive ratio of the optimal mechanism significantly unless if θ is relatively large. We will construct a family of Bayesian instances in which the competitive ratio of the optimal mechanism will increase by at most $\theta/1.5$ if the maximum cost is revealed to the buyer. Consider the impossibility example that we discuss in Section 7. Note that as n approaches ∞, the maximum bid quickly approaches $1-1/e$ from below. From our impossibility instance, construct an auxiliary instance by adding a seller s' with cost $1-1/e$ to the original instance. Also, let the budget in the auxiliary instance be $B' = B + 1 - 1/e$, where recall that B is the budget in the original instance. Furthermore, suppose that in the auxiliary instance, we reveal the cost of s' to the mechanism designer (i.e. we announce that s' has cost $1-1/e$). It can be shown that if the optimal mechanism in the original instance has competitive ratio α, then the optimal mechanism in the auxiliary instance would have competitive ratio $\frac{\alpha n+1}{n+1}$. The difference between the competitive ratios of the optimal mechanisms in the original and auxiliary instances would then be $\frac{\alpha n+1-\alpha(n+1)}{n+1} = \frac{1-\alpha}{n+1}$. Note that, by Chen et al. (2011), we have $\alpha \geq 1/3$ for any instance in the prior-free setting. This proves the claim.

Therefore, the benefit from revealing the maximum bid decreases quickly as the market becomes larger. For instance, for $n=20$ and $n=40$ bidders (which correspond to $\theta = 1/n$ in this example), this benefit would be around 0.02 and 0.01, respectively.

A.2. An Example for Mechanism Envy-Free

Suppose the buyer has budget $B = 13/3$, $S = \{s_1, s_2\}$ and sellers s_1, s_2 each owns a divisible item with costs 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 $f(d) = 1-d$ 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 8). 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 exceeds 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 9)

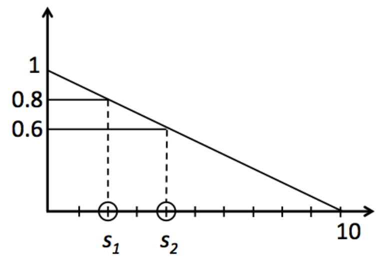

Figure 8 The mechanism is not budget feasible at $r = 10$.

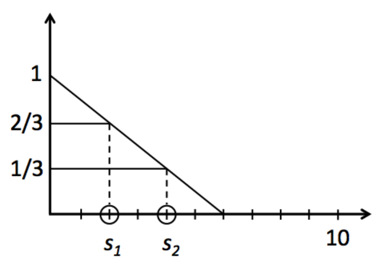

Figure 9 The mechanism is budget feasible at $r = 6$. Dotted lines represent the allocation from each seller.

A.3. Mechanism Envy-Free(f)

We show that Mechanism Envy-Free(f) is in fact envy-free, in the following sense.

Definition 6. A mechanism is envy-free if for any seller i who reports her true cost and for any other seller j we have:

$$P_i(\mathbf{c}) - c_i \cdot A_i(\mathbf{c}) \geq \frac{u_i}{u_j} \cdot (P_j(\mathbf{c}) - c_i \cdot A_j(\mathbf{c})) \tag{5}$$

.

Note that in the right-hand side

Proposition 3. *If seller i reports her true cost to Mechanism Envy-Free(f), then she would not envy any other seller in S.*

Proof. When the mechanism stops, it has chosen a monotone allocation rule f_{r^*} and the corresponding payment rules $\{P_{i,r^*}\}_{i\in S}$. Now, as a thought experiment, ignore the budget B and assume that we do not run the mechanism; instead, we use f_{r^*} and $\{P_{i,r^*}\}_{i\in S}$ directly to make the allocation and payments. By Myerson's characterization no seller has incentive to misreport her cost, which means she would not envy any other seller. □

Appendix B: Proofs from Section 6

First, we present the proof for Proposition 1, then the proof for Theorem 4, and after that, in Section B.3, we present the remaining proofs from Section 6.1.

B.1. Proof of Proposition 1

Recall the definition of $r_1, \ldots, r_n$ from Mechanism Truthful(f). Let r be the scaling ratio of Envy-Free(f). First, note that $r \le r_i$, for any $i \in S$. This holds simply because any payment rule $P(c)$ is decreasing in c. Therefore, under any fixed payment rule, the payment to seller i increases when her cost is set to 0. This implies that Envy-Free(f) exhausts the budget at a lower scaling ratio when c_i is set to 0, which means $r_i \le r$. Consequently, the buyer attains a lower level of utility under Truthful(f) than under Envy-Free(f).

It remains to show that the competitive ratio of Truthful(f) is asymptotically equal to the competitive ratio of Envy-Free(f) when f is log-concave. In Section F, we (independently) show this for $f = f^*$. The proof in here is essentially the same proof but with small adjustments for handling the more general case of a log-concave f. We start the analysis with the following two lemmas.

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

Proof. We need to prove that for $B_1, B_2 \geq 0$ and $0 \leq \lambda \leq 1$

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

Let x_i be the amount of item i we allocate to achieve $u^\star(B_1)$ and let y_i be the amount of item i we allocate to achieve $u^\star(B_2)$.

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

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

The cost paid by these allocations is simply

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

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 $\lambda u^\star(B_1) + (1-\lambda)u^\star(B_2)$. This proves that

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

□

Lemma 7. *For each $k \in \{1, \ldots, n\}$, $r_k \geq (1-\theta)r^*$.*

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

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

Here we used the fact that Q_{r^*} is a decreasing function. This implies that $\sum_{i=1}^{n} u_i Q_{(1-\theta)r^*}(\frac{c_i}{u_i}) \leq (1-\theta)B$. 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

$$u_k Q_{(1-\theta)r^*}(0) + \sum_{i \neq k} u_i Q_{(1-\theta)r^*}(\frac{u_i}{c_i}) \leq (1-\theta)B + u_k \left(Q_{(1-\theta)r^*}(0) - Q_{(1-\theta)r^*}(\frac{c_k}{u_k}) \right). \qquad (6)$$

Note that $Q_{(1-\theta)r^*}(\cdot)$ is defined as the area of the shaded region as seen in Figure 1. Therefore one can crudely upper bound the difference $Q_{(1-\theta)r^*}(0) - Q_{(1-\theta)r^*}(d)$ by $d \times f_{(1-\theta)r^*}(0)$ for any $d \geq 0$. Now letting $d = \frac{c_k}{u_k}$, and substituting in (6) we get

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

$$= (1-\theta)B + c_k \leq B.$$

This completes the proof. □

We now can finish the proof for Proposition 1. Let ζ', ζ respectively denote the competitive ratios of mechanisms Truthful(f) and Envy-Free(f). We will show that $\zeta' \to \zeta$ as $\theta \to 0$. 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 $\{1, \ldots, n\}$ 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

$$B = u_k Q_{r_k}(0) + \sum_{i \neq k} u_i Q_{r_k}(\frac{c_i}{u_i}) \leq u_k \cdot (e-1) + \sum_{i \neq k} u_i Q_{r_k}(\frac{c_i}{u_i}). \tag{7}$$

Here we used the fact that $Q_{r_k}(0) \leq Q_{r^*}(0) \leq e-1$ (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 (7) we get

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

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

$$\sum_{i=1}^n u_i f_{r^+}(\frac{c_i}{u_i}) \geq (1 - \theta \cdot \frac{e-1}{s}) \zeta \cdot u^\star(B). \tag{8}$$

For an item $i \in \mathcal{I}$, we have $r_i \geq (1-\theta)r^* = 1-\theta$ (using Lemma 7). Therefore

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

Since f is log-concave, then $\frac{f(\frac{1}{r_i}x)}{f(x)}$ 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 $\alpha = \frac{f(\frac{s}{r_i})}{f(s)}$, 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 $i \in \{1, \ldots, n\}$, $f_{r_i}(\frac{c_i}{u_i}) \geq \alpha f_{r^+}(\frac{c_i}{u_i})$. Combining this with (8) we get

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

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

$$\alpha(1 - \theta \cdot \frac{e-1}{s}) \cdot \zeta \tag{9}$$

For any fixed s, strictly smaller than $e-1$, one can observe that (9) approaches ζ as $\theta \to 0$. This completes the proof for Proposition 1. We do not attempt to optimize over the choice of s for brevity.

B.2. Proof of Theorem 4

We will use the following lemmas in the proof for Theorem 4.

Lemma 8. *$\widehat{f}(x)$ is strictly increasing in x.*

Proof. $\widehat{f}(f(c)) = Q_f(c) - c$, where $Q_f(c)$ and $f(c)$ are respectively strictly decreasing in c. This implies the claim. □

In the next two lemmas we show that Scaling f with ratio r along the horizontal axis (i.e. transforming it to f_r) corresponds to scaling $\widehat{f}$ along the vertical axis with ratio r.

LEMMA 9. *For any standard allocation rule f, we have $\widehat{f_r}(x) = r\widehat{f}(x)$.*

Proof.

$$\widehat{f_r}(x) = Q_{f_r}(x) - {f_r}^{-1}(x)$$
$$= rQ_f(x) - rf^{-1}(x) = r\widehat{f}(x).$$

□

LEMMA 10. *For any $r > 0$,*

$$\inf\left\{x : (x,0) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}\right\} = \inf\left\{x : (x,0) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f_r}\right\}\right\}.$$

Proof. By Lemma 9, scaling f with ratio r along the horizontal axis (i.e. transforming it to f_r) corresponds to scaling $\widehat{f}$ along the vertical axis with ratio r. Scaling $\widehat{f}$ along the vertical axis with ratio r also scales its lower convex hull with ratio r. Formally,

$$(x,y) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\} \Longleftrightarrow (x, ry) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f_r}\right\}.$$

Therefore, $(\alpha, 0) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}$ implies $(\alpha, 0) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f_r}\right\}$, for any α. This proves the claim. □

LEMMA 11. *For any standard allocation rule f, the competitive ratio of Envy-Free(f) is at most*

$$\inf\left\{x : (x,0) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}\right\}.$$

Proof. Let c_0 be the cost for which $f(c_0) = \inf\left\{x : (x,0) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}\right\}$. There exists c_1, c_2 such that $c_1 \leq c_0 \leq c_2$ and $f(c_0) = \gamma_1 f(c_1) + \gamma_2 f(c_2)$ for some $\gamma_1, \gamma_2 \geq 0$ with $\gamma_1 + \gamma_2 = 1$. (See Figure 10) We construct an instance with n sellers: $\gamma_1 n$ of them have cost c_1 and $\gamma_2 n$ of them have cost c_2. As you will see in a moment, we will approach n to infinity; this will solve the issue of the possible non-integrality of $\gamma_1 n, \gamma_2 n$. So, without getting into the tedious details, we just remark that it is safe to assume $\gamma_1 n, \gamma_2 n$ are integers. Define the budget of the buyer to be $B = n\gamma_1 c_1 + n\gamma_2 c_2$.

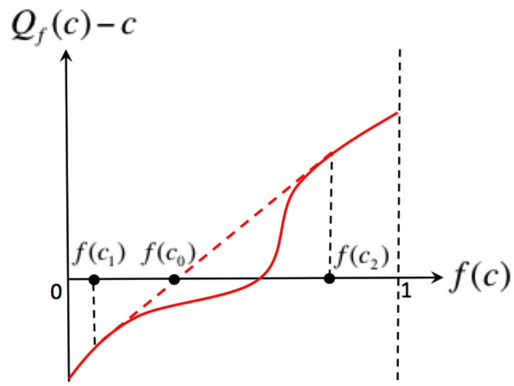

Figure 10 The solid red curve represents $\widehat{f}$ and the dotted red line represents part of the boundary of $\mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}$.

By the definition of c_1, c_2, we have

$$\begin{aligned}0 &= \gamma_1 \widehat{f}(c_1) + \gamma_2 \widehat{f}(c_2) \\ &= \gamma_1 \left(Q_f(c_1) - c_1\right) + \gamma_2 \left(Q_f(c_2) - c_2\right) \\ &\Rightarrow \gamma_1 Q_f(c_1) + \gamma_2 Q_f(c_2) = B\end{aligned}$$

Consequently, Envy-Free(f) exhausts the budget at $r = 1$. This implies that the allocation of Envy-Free(f) will be $n(\gamma_1 f(c_1) + \gamma_2 f(c_2))$. In the other hand, note that $U^* = n$, since the buyer's budget is equal to the sum of the costs of all items. Therefore, the competitive ratio of Envy-Free(f) is $\gamma_1 f(c_1) + \gamma_2 f(c_2) = f(c_0)$ for the defined instance. □

Proof of Theorem 4: We define some notation first. Let $\mathrm{supp}(f) = [0, \overline{c}]$. Because of Lemma 10, WLOG we can assume that the scaling ratio in Envy-Free(f) is $r = 1$.

Define $\alpha = \inf\left\{x : (x, 0) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}\right\}$. By the separating hyper plane theorem, there should be a line l passing through $(\alpha, 0)$ with all the points in $\mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}$ falling on the same side of this line. Let $y = sx + b$ denote the line equation for l. We can show that $0 < s < \infty$

Claim 3. *The slope s is a strictly positive number, i.e. $0 < s < \infty$.*

The proof for this claim comes after the proof of the theorem.

Now, suppose that Envy-Free(f) allocates $x_i = f(c_i)$ units from seller i. Then, we should have

$$\widehat{f}(x_i) \leq sx_i + b, \tag{10}$$

since $(x_i, \widehat{f}(x_i)) \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}$. Summing up this inequality for all i implies that

$$\sum_{i=1}^{n} \widehat{f}(x_i) = \sum_{i=1}^{n} Q_f(c_i) - c_i \leq \sum_{i=1}^{n} sx_i + bn. \tag{11}$$

We are done if we have $\sum_{i=1}^{n} c_i \leq B$. In this case, the above inequality would imply that $0 \leq \sum_{i=1}^{n} sx_i + bn$, or equivalently, $\frac{-bn}{s} \leq \sum_{i=1}^{n} x_i$. Now, note that on the left-hand side, $n = U^*$ and $\frac{-b}{s}$ is exactly the x-intercept of l. In the other hand, the right-hand side is just the total utility attained by the mechanism. So, the theorem is proved when $\sum_{i=1}^{n} c_i \leq B$ holds. It remains to address the case of $\sum_{i=1}^{n} c_i > B$.

We write a stronger version of (10). For any $\zeta \in [0,1]$ and $c \in [0, \overline{c}]$, we will show that

$$Q_f(c_i) - \zeta c_i \leq sx_i + \zeta b. \tag{12}$$

Note that (12) coincides with (10) when $\zeta = 1$. In the next claim, we show that (12) holds also when $\zeta = 0$. The proof comes after the proof of the theorem.

Claim 4. *For all $c \in [0, \overline{c}]$, $Q_f(c) \leq sf(c)$.*

Since (12) holds for $\zeta \in \{0, 1\}$, then it also holds for any $\zeta \in [0,1]$. (See this by taking a convex combination of the inequalities written for $\zeta = 0$ and $\zeta = 1$) Suppose that the omniscient mechanism (which picks the optimal solution of the underlying optimization problem) allocates a fraction ζ_i from the item of seller $i \in S$. Then, we can use (12) to write

$$Q_f(c_i) - \zeta_i c_i \leq sf(c_i) + \zeta_i b. \tag{13}$$

We prove the lemma by adding up (13) for all $i \in S$: On the left-hand side, we get $\sum_{i \in S} Q_f(c_i) - \zeta_i c_i$, which is equal to 0, since $\sum_{i \in S} Q_f(c_i) = B = \sum_{i \in S} \zeta_i c_i$. So, adding up these inequalities implies

$$\begin{aligned} &0 \leq \sum_{i \in S} sf(c_i) + \zeta_i b, \\ &\Rightarrow \frac{-b}{s} \cdot \sum_{i \in S} \zeta_i \leq \sum_{i \in S} f(c_i). \end{aligned} \tag{14}$$

The left-hand side of (14) is precisely α times the utility of the omniscient mechanism, and the right-hand side is the allocation of Envy-Free(f). So, the competitive ratio of Envy-Free(f) is at least α. The proof is completed by Lemma 11, which shows that the competitive ratio is at most α. □

Proof of Claim 3: First, we show that $\alpha > 0$. $\widehat{f}$ intersects the horizontal axis at a strictly positive coordinate: this holds since $\widehat{f}(0) < 0$, and since $\widehat{f}$ is strictly increasing, by Lemma 8. Therefore $\mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}$ should also intersect the horizontal axis at a strictly positive coordinate, which means $\alpha > 0$.

Having $\alpha > 0$ rules out the possibility of $s = 0$ and $s = \infty$ (we leave out the details). Therefore, it remains to show that s is not negative. Consider the points $p = (f(0), \widehat{f}(f(0)))$ $p' = (f(\overline{c}), \widehat{f}(f(\overline{c})))$ (marked in Figure 11). Since $p, p' \in \mathsf{Conv}^{\downarrow}\left\{\widehat{f}\right\}$, then they should be on the same side of l. In the

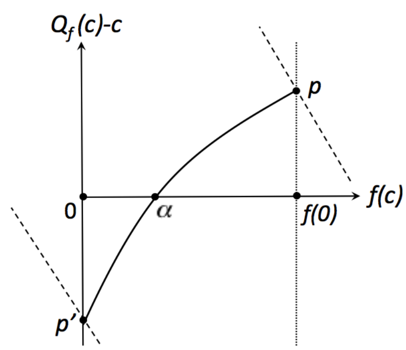

Figure 11 A line with a negative slope can not pass through $(\alpha, 0)$ if it has p, p' on the same side.

other hand, $\widehat{f}(f(\overline{c})) < 0$, $\widehat{f}(f(0)) > 0$, and l should pass through $(\alpha, 0)$. Therefore, the slope of l cannot be negative (see Figure 11 for a proof by picture). □

Proof of Claim 4: Let $q = (f(c), \widehat{f}(f(c)))$. Also, let $q' = (0, b)$ denote the point at which l crosses the vertical axis (note that $q' = p'$, where p' was defined in the proof of Claim 3). Define l' to be line passing through q, q'. The slope of l', namely s', should be smaller than the slope of l, by the definition of l. Therefore, we can write

$$s' = \frac{Q_f(c) - c - (Q_f(\overline{c}) - \overline{c})}{f(c) - f(\overline{c})} \leq s,$$

$$\Rightarrow \frac{Q_f(c) - Q_f(\bar{c})}{f(c) - f(\bar{c})} = \frac{Q_f(c)}{f(c)} \leq s.$$

□

B.3. Proofs from Section 6.1

Proof of Proposition 2: Let $f = f^*$. First, we compute the closed form expression for $\widehat{f}(c)$.

$$Q_f(c) = cf(c) + (e-c)\ln(e-c) + c - (e-1),$$
$$\widehat{f}(c) = Q_f(c) - c$$

Now, it is straight-forward to verify that $\frac{\partial \widehat{f}(c)}{\partial f(c)} = e$; this would imply that $\widehat{f}$ is a straight line. (This is observable in Figure 3(a), where we have plotted $\widehat{f}$) We also have to verify that the x-intercept of this line is $1 - 1/e$: For $c_0 = e - e^{1-1/e}$, we have $f(c_0) = 1 - 1/e$ and $\widehat{f}(f(c_0)) = 0$.

It remains to show that $\widehat{f}_r$ is a straight line for any $r > 0$. This is implied by Lemma 9 which says $\widehat{f}_r(x) = r\widehat{f}(x)$. The proof is complete since scaling a straight line along the vertical axis produces a straight line. □

Proof of Lemma 2: The proof is by contradiction. Suppose

$$\widehat{f} \cap \widehat{g} = \emptyset. \tag{15}$$

Q_f, Q_g respectively denote the payment rules corresponding to allocation rules f, g. Without loss of generality, suppose $Q_f(0) < Q_g(0)$. This means

$$\widehat{f}(x) < \widehat{g}(x), \quad \forall x \in \mathrm{supp}(\widehat{f}) \cap \mathrm{supp}(\widehat{g}). \tag{16}$$

Let $[0, c_f] = \mathrm{supp}(\widehat{f})$ and $[0, c_g] = \mathrm{supp}(\widehat{g})$. Now, (15) and (16) imply that $c_f > c_g$. Since $Q_f(0) < Q_g(0)$, the latter fact implies that there exists a number $x^* \in [0, c_g]$ for which $f(x^*) = g(x^*)$. Choose x^* to be the largest possible such number. This would imply $Q_f(x^*) > Q_g(x^*)$, which means $\widehat{f}(x^*) > \widehat{g}(x^*)$. Contradiction. □

Proof of Lemma 3: The proof is by contradiction; suppose the allocation rule $g \neq f^*$ has competitive ratio $1 - 1/e$. Let Q_g denote the unit-payment rule corresponding to allocation rule g.

Also, let $[0,c_g]=\mathrm{supp}(\widehat{g})$. By Lemma 2 and Proposition 2, there must exist a line l passing through $(1-1/e,0)$ which has at least two points in common with $\widehat{g}$ (Recall the proof sketch for Lemma 3 that was presented in Section 6.1). Figure 12 demonstrates such a line.

By Proposition 2, there exists $r>0$ such that l corresponds to $\widehat{f_r^*}$. By Lemma 10, without loss of generality we can assume that $r=1$ (this is just a scaling). For notational simplicity, we will denote f_r^* by $f(=f^*)$ from now on. Let Q_f denote the unit-payment rule corresponding to allocation rule f. Also, let $[0,c_f]=\mathrm{supp}(\widehat{f})$. We consider two cases to prove the theorem: $c_f<c_g$ and $c_f=c_g$. (Note that since $l=\widehat{f}$ lies above $\widehat{g}$, then $\widehat{f}(0)\geq\widehat{g}(0)$ holds, which rules out the case of $c_f>c_g$.)

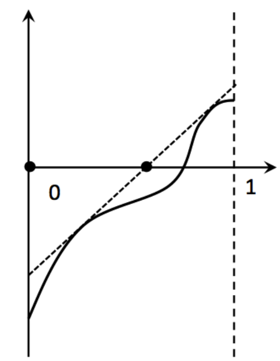

Figure 12 The dotted line and solid curve respectively represents $l=\widehat{f}$ and $\widehat{g}$.

Case $c_f<c_g$. In this case, there must exists some $c^*<c_f$ such that $g(c^*)=f(c^*)$ and $g(c)>f(c)$ for all $c>c^*$. Let $x^*=f(c^*)$. We then should have $Q_g(c^*)>Q_f(c^*)$, and therefore, $\widehat{g}(x^*)>\widehat{f}(x^*)$. But this contradicts the definition of l.

Case $c_f=c_g$. In this case, we must have $\widehat{f}(x)\geq\widehat{g}(x)$ for all $x\in[0,1]$, otherwise, $(1-1/e,0)$ would be an interior point of $\mathsf{Conv}^{\downarrow}\{\widehat{g}\}$, which is a contradiction. Define

$$\overline{c}=\sup_{c\leq c_f}\{c:f(c)\neq g(c)\},$$

$$\underline{c}=\sup_{c\leq c_f}\{c:f(c)>g(c)\}$$

First, we consider the case $\underline{c}<\overline{c}$. In this case, we will have $Q_g(\underline{c})>Q_f(\underline{c})$. This fact and the fact that $f(\underline{c})=g(\underline{c})$ imply that $\widehat{g}(\underline{c})>\widehat{f}(\underline{c})$, which is a contradiction. So, we can assume $\underline{c}=\overline{c}$.

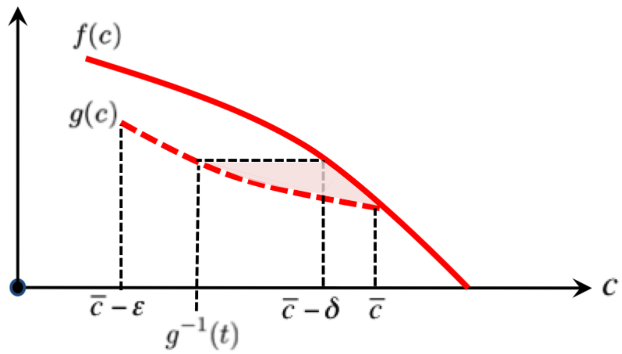

Figure 13 Choosing a sufficiently small δ.

By the continuity of f, g, there must exist $\epsilon > 0$ such that $f(c) > g(c)$ for all $c \in (\overline{c} - \epsilon, \overline{c})$. We will show that there exists a positive $\delta < \epsilon$ such that $\widehat{g}(f(\overline{c} - \delta)) > \widehat{f}(f(\overline{c} - \delta))$, which would be a contradiction.

Let $t = f(\overline{c} - \delta)$. By the continuity of f, g, for a sufficiently small choice of $\delta > 0$, we would have $g^{-1}(t) \in (\overline{c} - \epsilon, \overline{c})$, as it is show in Figure 13. Let $L = f^{-1}(t) - g^{-1}(t)$.

Claim 5. *Suppose $h : [0, a] \to \mathbb{R}$ is a continues function such that $h(0) = 0$ and $h(b) > 0$ for any $b \in (0, a]$. Then, there exists $\epsilon > 0$ such that $\epsilon > \int_0^\epsilon h(x)\,\mathrm{d}x$.*

This claim is proved after this proof. Applying this claim on $f^{-1} - g^{-1}$ would imply that $\delta > 0$ could be chosen sufficiently small so that

$$L > Q_f(f^{-1}(t)) - Q_g(g^{-1}(t)). \tag{17}$$

Note that the RHS is the shaded area in Figure 13. In the other hand, note that

$$g^{-1}(t) = f^{-1}(t) - L. \tag{18}$$

(17) and (18) together imply that $\widehat{g}(t) > \widehat{f}(t)$. Contradiction. □

Proof of Claim 5: For any positive $\epsilon \leq h(a)$ define

$$g(\epsilon) = \inf_{x \in [0,a]} \{x : h(x) > \epsilon\}.$$

Note that for any number $z > 0$, if $g(z) > z$, then $\int_0^z h(x)\,\mathrm{d}x < z^2$.

Let $\delta = \min\{a, 1/2\}$. If $g(\delta) > \delta$, then $\int_0^\delta h(x)\,\mathrm{d}\,x < \delta^2 < \delta$; therefore, the lemma will be proved by setting $\epsilon = \delta$. If $g(\delta) \le \delta$, then set $\epsilon = g(\delta)$; then, we would have

$$\int_0^\epsilon h(x)\,\mathrm{d}\,x < \delta\epsilon < \epsilon.$$

This proves the claim. □

Proof of Lemma 1: Proposition 2 says that $\widehat{f^*}$ is a straight line with x-intercept $1-1/e$. Together with Theorem 4 they imply that Mechanism Envy-Free(f^*) has competitive ratio $1-1/e$. In the other hand, Lemma 3 shows that the competitive ratio of Envy-Free(f) is strictly less than $1-1/e$ for any standard allocation rule $f \neq f^*$. This completes the proof. □

Appendix C: Theorem 1

We present the complete proof for Theorem 1. The proof is self-contained. It is presented in this section and sections D, E, and F.

Proof of Theorem 1: In Lemma 13 (this section), we show that Mechanism Truthful(f) is individually rational, truthful, and budget-feasible for any given standard allocation rule f. In Lemma 15 (Section D), we show that mechanism Envy-Free(f) has competitive ratio $1-1/e$ for $f = f^*$, and finally, in Lemma 19 (Section F) we show that Truthful(f) has competitive ratio $1-1/e$. This result, together with Theorem 3 imply that $f = f^*$ is an optimal allocation rule for Truthful(f).

□

Here, we will prove that Mechanism Truthful(f) is individually rational, truthful, and budget-feasible for any given standard allocation rule f. First, we state the following lemma, which comes handy in the proof.

Lemma 12. *For any seller $i \in S$ we have $r^* \ge r_i$.*

Proof. The proof is based on the fact that $P_{i,r}(x)$ is an increasing function of r (for a fixed x) and 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

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

where the first inequality holds because $P_{j,r^*}(x)$ is a decreasing function of x and the second inequality holds because $r^* < r_i$. However, note that the above inequalities imply $B < \sum_{j \in S} P_{j,r_i}(c'_j)$, 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 $c'_1, \ldots, c'_n$, and so it can not be larger than B. □

LEMMA 13. *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 characterization of truthful mechanisms (Myerson (1981)) 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 why $p_i \leq p'_i$ holds, 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 12, it must be the case that $P_{i,r_i}(c_i) \leq P_{i,r^*}(c_i)$. □

Appendix D: Theorem 1: Finding an optimal f for the Mechanism Envy-Free(f)

Here, we prove that Mechanism Envy-Free(f) has competitive ratio $1 - 1/e$ for $f(x) = \ln(e - x)$. Although Mechanism Envy-Free(f) is not truthful, its analysis forms the core of our analysis for the truthful mechanism. In this section, we analyze Mechanism Envy-Free(f) assuming that the true costs are known. We defer the analysis for Mechanism Truthful(f) to Section F, where we prove that its competitive ratio is (asymptotically) equal to the competitive ratio of Mechanism Envy-Free(f), for the same choice of f.

D.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 $\frac{c_i}{u_i}$, is in terms of $f_r(\frac{c_i}{u_i})$. This alternative definition is denoted by G_r, and is defined such that $Q_r(\frac{c_i}{u_i}) = G_r(f(\frac{c_i}{u_i}))$. 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. Formally, for $y = f_r(\frac{c_i}{u_i})$ we define

$$G_r(y) = \int_0^y g_r(x)\,\mathrm{d}\,x = Q_r(\frac{c_i}{u_i}).$$

We also denote g_1 and G_1 respectively by g and G. These functions are computed below for the optimal allocation rule. The proof is straight-forward and is relegated to Section A.3.

Proposition 4. *For the standard allocation rule $f(x) = \ln(e - x)$, we have $g(x) = e - e^x$ and $G(x) = ex - e^x + 1$. Also, we have that $f_r(x) = \ln\left(\frac{er-x}{r}\right)$.*

Proof. For the standard allocation rule $f(x) = \ln(e - x)$, we need to show that $g(x) = e - e^x$ and $G(x) = ex - e^x + 1$. Note that by definition, $g(x)$ is the inverse of $f(x) = \ln(e - x)$; a straight forward calculation shows that $f^{-1}(x) = e - e^x$. Also, by definition, we have:

$$G(x) = \int_0^x g(y)\,\mathrm{d}\,y = ex - e^x + 1.$$

□

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 14. *For any x, α such that $0 \le x, \alpha \le 1$ we have*

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

Proof.

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

□

D.2. Competitive Ratio of Mechanism Envy-Free(f)

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

Lemma 15. *If sellers report their true costs, then Mechanism Envy-Free(f) has competitive ratio $1 - 1/e$.*

Proof. Observe that w.l.o.g. we can assume $r^* = 1$: If $r^* \neq 1$, then we can transform the given instance to a new instance for which the mechanism has scaling ratio 1. More precisely, there exists some $\beta > 0$ such that if we multiply the budget and the reported costs by β, the scaling ratio of the mechanism becomes equal to 1 (intuitively, this operation can be seen as a contraction along the horizontal axis). 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 14 to write the following set of inequalities:

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

where α_i is the fraction that is allocated from seller i in the optimal solution (recall that we are 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:

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

By adding up these inequalities, we get:

$$\sum_{i \in S} u_i \cdot (G(x_i) - \alpha_i \cdot g(x_i)) \le e \cdot \sum_{i \in S} u_i \cdot (x_i - \alpha_i \cdot (1 - 1/e)) \,. \tag{19}$$

Now, we show that if

$$0 \le \sum_{i \in S} u_i \cdot (G(x_i) - \alpha_i \cdot g(x_i)) \,, \tag{20}$$

then the lemma is proved using (19) and (20). First we show why (19) and (20) prove the lemma, and then in the end, we prove (20) itself.

Observe that (19) and (20) imply that

$$0 \le \sum_{i \in S} u_i \cdot (x_i - \alpha_i \cdot (1 - 1/e)) \,. \tag{21}$$

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 (21) implies

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

This would prove the lemma.

So, it only remains to show that (20) holds: First observe that $\sum_{i \in S} u_i \cdot G(x_i) = B$, since the sum represents the total payment of Envy-Free(f). Also, see that $\sum_{i \in S} \alpha_i u_i \cdot g(x_i) \le B$, since this sum is a lower bound on the cost of the optimal solution, which is at most B. □

The analysis for our truthful mechanism is presented in Sections E and F, where we prove that its competitive ratio is (asymptotically) equal to the competitive ratio of Mechanism Envy-Free(f), for the same choice of f. In Section E, we analyze the truthful mechanism for the case of unit utilities, and then we analyze the case of general utilities in Section F.

Appendix E: Theorem 1: Competitive ratio of Truthful(f) for unit utilities

In this section we focus on the special case when all the utilities are equal to 1, and prove that Mechanism Truthful(f) has competitive ratio $1 - 1/e$ for $f = f^*$. The proof for the case of general utilities is more intricate and appears in Section F.

First, we state a simple lemma that comes handy in the analysis. WLOG, assume that $c_1 \le c_2 \le \ldots \le c_n$. (Note that the assumption of unit utilities implies that $c_i/u_i = c_i$ for any seller i)

LEMMA 16. $r_1 \geq r_2 \geq \ldots \geq r_n$.

Proof. 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:

$$Q_{r_i}(0) + \sum_{k \in S \setminus \{i\}} Q_{r_i}(c_k) = B. \tag{22}$$

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

$$Q_{r_j}(0) + \sum_{k \in S \setminus \{j\}} Q_{r_j}(c_k) = B. \tag{23}$$

Taking the difference between (22) and (23) implies that

$$\Big(Q_{r_i}(0) - Q_{r_j}(0)\Big) + \Big(Q_{r_i}(c_j) - Q_{r_j}(c_i)\Big) = 0 \tag{24}$$

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 (24). □

We are now ready to prove the main claim of this section.

LEMMA 17. *Mechanism Truthful(f) has competitive ratio at least $1 - 1/e$ when all the items have utility equal to* 1.

Proof. Let $U^* = u^\star(B)$ and 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 $f_{r_n}(c_i) \leq f_{r_i}(c_i)$, note that $r_n \leq r_i$ due to Lemma 16 which implies $f_{r_n}(c_i) \leq f_{r_i}(c_i)$)

The proof proceeds by providing a lower bound on $\sum_{i \in S} f_{r_n}(c_i)$ in terms of U^*. 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

$$U_n \geq (1-1/e) \cdot (1-o(1)) \cdot U^*. \tag{25}$$

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 scaling ratio 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 15 (or Lemma 1) on the auxiliary instance, we have $U_n \geq (1-1/e) \cdot U^*_{\mathsf{aux}}$. So, if we show that

$$U^*_{\mathsf{aux}} \geq (1-o(1)) \cdot U^* \tag{26}$$

then (25) holds and the proof is complete.

We use Lemma 6 to prove (26): First, we show that $B' \geq (1-o(1)) \cdot B$; then, applying Lemma 6 would imply that $u^\star(B') \geq (1-o(1)) \cdot u^\star(B)$, which is identical to (26) by definition. So, the proof is complete if we show that $B' \geq (1-o(1)) \cdot B$.

We prove that $B' \geq (1-\alpha \cdot \frac{c_{\mathsf{max}}}{B}) \cdot B$, where α is a constant with value $(e-e^{1-1/e})^{-1} \approx 6/5$. This would prove the Lemma due to the Small Bidders assumption (i.e. $\frac{c_{\mathsf{max}}}{B} \to 0$). First, observe that

$$\begin{aligned} B = Q_{r_n}(0) + \sum_{i \in S \setminus \{n\}} Q_{r_n}(c_i) \leq Q_{r_n}(0) + B' \\ \Rightarrow B' \geq B - Q_{r_n}(0) \geq B - r_n. \end{aligned} \tag{27}$$

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

$$\begin{aligned} B \geq \frac{\overline{c}}{c_{\mathsf{max}}} \cdot B = r_n(e-e^{1-1/e}) \cdot \frac{B}{c_{\mathsf{max}}} \\ \Rightarrow c_{\mathsf{max}} \cdot (e-e^{1-1/e})^{-1} \geq r_n. \end{aligned} \tag{28}$$

Combining (27) and (28) implies $B' \geq (1-\alpha \cdot \frac{c_{\mathsf{max}}}{B}) \cdot B$ with the promised value for α. $\square$

Proposition 5 **(Corollary of Lemma 17)**. *When all items have utility* 1*, the competitive ratios of Truthful*(f) *and Envy-Free*(f) *are equal.*

Proof. Let ζ', ζ respectively denote the competitive ratios of mechanisms Truthful(f) and Envy-Free(f). We will show that $\zeta' \to \zeta$ as $\theta \to 0$. The proof is essentially the same proof we presented for Lemma 17, except that we replace the competitive ratio $1-1/e$ with ζ, and instead of picking $\bar{c}$ so that $f_{r_n}(\bar{c}) = 1-1/e$, we pick $\bar{c}$ so that $f_{r_n}(\bar{c}) = \zeta$. The existence of such $\bar{c}$ is guaranteed because f is a standard allocation rule. The rest of the proof remains the same. □

Proposition 6 **(Corollary of Lemma 17)**. *Define* $\theta = \max_{i\in S}\{\frac{c_i u_i}{B}\}$*. Then, as* $\theta \to 0$*, the competitive ratios of Truthful*(f) *approaches the competitive ratio of Envy-Free*(f)*.*

Proof. The proof is essentially the proof of Lemma 17 with small adjustments. Roughly speaking, the condition on θ ensure that, similar to the case of unit utilities, the utilities are sufficiently small. Therefore, the scaling ratio does not change a lot when the cost of a seller is set to 0. We omit the details. □

Appendix F: Theorem 1: Competitive ratio of Truthful(f) for general utilities

In this section we will prove that the competitive ratio of mechanism Truthful(f) approaches $1-1/e$ as θ, the market's bid-budget ratio, approaches 0. We emphasize that here we dismiss the extra assumption that was made in Section E: 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 18. *For each* $k \in \{1,\ldots,n\}$*,* $r_k \geq (1-\theta)r^\star$*.*

Proof. We just need to prove that $f_{(1-\theta)r^\star}$ 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

$$Q_{(1-\theta)r^\star}(x) = (1-\theta)Q_{r^\star}(\frac{x}{1-\theta}) \leq (1-\theta)Q_{r^\star}(x).$$

Here we used the fact that $Q_{r^\star}$ is a decreasing function. This implies that $\sum_{i=1}^n u_i Q_{(1-\theta)r^\star}(\frac{c_i}{u_i}) \leq (1-\theta)B$. This expression is the budget consumed by the rule $f_{(1-\theta)r^\star}$ 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

$$u_k Q_{(1-\theta)r^\star}(0) + \sum_{i \neq k} u_i Q_{(1-\theta)r^\star}(\frac{u_i}{c_i}) \leq (1-\theta)B + u_k \left(Q_{(1-\theta)r^\star}(0) - Q_{(1-\theta)r^\star}(\frac{c_k}{u_k}) \right). \tag{29}$$

Note that $Q_{(1-\theta)r^\star}(\cdot)$ is defined as the area of the shaded region as seen in figure 1. Therefore one can crudely upper bound the difference $Q_{(1-\theta)r^\star}(0) - Q_{(1-\theta)r^\star}(x)$ by $x \times f_{(1-\theta)r^\star}(0)$ for any $x \geq 0$. Now letting $x = \frac{c_k}{u_k}$, and substituting in inequality 29 we get

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

This completes the proof. □

Lemma 19. *Mechanism Truthful(f) has a competitive ratio approaching* $1 - \frac{1}{e}$ *as* θ *approaches* 0.

Proof. W.l.o.g. assume that $r^\star = 1$ (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 $\{1, \ldots, n\}$ 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 $r^+ = r^\star = 1$. 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 $r^+ = r^\star$, 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

$$B = u_k Q_{r_k}(0) + \sum_{i \neq k} u_i Q_{r_k}(\frac{c_i}{u_i}) \leq u_k + \sum_{i \neq k} u_i Q_{r_k}(\frac{c_i}{u_i}). \tag{30}$$

Here we used the fact that $Q_{r_k}(0) \leq Q_{r^\star}(0) = 1$ (since we assumed $r^\star = 1$). Note that $B' \geq \sum_{i \neq k} u_i Q_{r^+}(\frac{c_i}{u_i})$. Combining this with the inequality 30 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 $u^\star(B') \geq (1-\frac{\theta}{s})u^\star(B)$. But we also know from lemma 15 that the utility achieved by f_{r^+} is at least $(1-\frac{1}{e})u^\star(B')$. Therefore we have

$$\sum_{i=1}^{n} u_i f_{r^+}(\frac{c_i}{u_i}) \geq (1-\frac{\theta}{s})(1-\frac{1}{e})u^\star(B) \tag{31}$$

For an item $i \in \mathcal{I}$, we have $r_i \geq (1-\theta)r^\star = 1-\theta$ (we used lemma 18). 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 $\alpha = \frac{f(\frac{s}{r_i})}{f(s)}$, 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 $i \in \{1,\ldots,n\}$, $f_{r_i}(\frac{c_i}{u_i}) \geq \alpha f_{r^+}(\frac{c_i}{u_i})$. Combining this with (31) we get

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

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

$$\alpha(1-\frac{\theta}{s})(1-\frac{1}{e}) = \frac{\ln(e-\frac{s}{1-\theta})}{\ln(e-s)}(1-\frac{\theta}{s})(1-\frac{1}{e}) \tag{32}$$

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

According to the above analysis, we can also compute a lower bound on the competitive ratio of the optimal mechanism for the general case.

COROLLARY 2. *Competitive ratio of Truthful(f) is at least* $\frac{f^*(1)}{f^*(1-\theta)} \cdot (1-\frac{\theta}{1-\theta}) \cdot (1-1/e)$.

Proof. For any positive $s < 1$, the right-hand side of (32) gives a lower bound on the competitive ratio. Without optimizing over this free parameter, we choose $s = 1-\theta$ to simplify the calculations, which leads to $\frac{f^*(1)}{f^*(1-\theta)} \cdot (1-\frac{\theta}{1-\theta}) \cdot (1-1/e)$, where recall that $f^*(x) = \ln(e-x)$. □

Appendix G: Mechanisms for Indivisible Items

In this section, we prove Theorem 6 by converting our mechanism for divisible items to a mechanism for indivisible items with the same competitive ratio. Let us call any valid allocation for the problem with divisible (indivisible) items a *fractional* (*integral*) allocation. The idea is running the mechanism for divisible items to obtain a fractional allocation, and then transforming it to an integral allocation using a randomized process. Since the final allocation would be a random variable itself, we redefine the incentive-compatibility constraint (1) as

$$\mathbb{E}\left[P_i(\bar{c}_i, \mathbf{c_{-i}}) - c_i \cdot A_i(\bar{c}_i, \mathbf{c_{-i}})\right] \leq \mathbb{E}\left[P_i(c_i, \mathbf{c_{-i}}) - c_i \cdot A_i(c_i, \mathbf{c_{-i}})\right],$$

where the expectation is taken over the integral allocations (and the corresponding payments) that could be produced by the mechanism.

We design a rounding process that takes the fractional allocation as its input and outputs an integral allocation with its associated payments. By the properties of our rounding process, the resulting mechanism would be individual rational, truthful, and budget-feasible; also, it would have competitive ratio $1 - 1/e$ under the Small Bidders assumption.

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 competitive ratio would remain $1 - 1/e$. First we explain these properties in Section G.1, then, we state our rounding procedure and prove that it satisfies these desired properties in Section G.2.

G.1. Properties of the Rounding Procedure

Let $\tilde{x}_1, \ldots, \tilde{x}_n$ represent a fractional allocation where $\tilde{x}_i$ denotes the allocated fraction from seller $i \in S$; also, let $\tilde{p}_1, \ldots, \tilde{p}_n$ be the associated payments for this allocation. We round this solution to an integral solution, represented by the allocation $x_1, \ldots, x_n$ and payments $p_1, \ldots, p_n$, 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. $\sum_{i \in S} p_i \leq B + c_{\mathsf{max}}$.

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}.[17] The theorem mentions that the resulting mechanism might bear an additive error of at most c_{max}. Note that c_{max} is supposed to be very small, since c_{max}/B goes to 0 by the Small Bidders assumption. However, we can still convert this mechanism to a strictly budget feasible mechanism if one is desired. More details about this appear in Section G.3.

G.2. Description of the Rounding Procedure

In this section, we focus in designing a rounding procedure which satisfies Properties 1, 2 and 3. Our procedure is a randomized rounding procedure (See **?** for a general survey on randomized rounding methods for designing algorithms). It takes a fractional allocation as its input and "implements" it as a convex combination of integral or "almost" integral allocations. First, we 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 20. *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 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 $\sum_{i \in S} y_i \cdot \frac{\tilde{p}_i}{\tilde{x}_i} \le B$. 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. □

Outline of the Rounding Procedure The procedure accepts the fractional allocation constructed by the mechanism, i.e. $\tilde{x} = (\tilde{x}_1, \ldots, \tilde{x}_n)$, 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 coefficient in the convex combination. Finally, 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 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.

Lemma 21. *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:

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

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 $p(x) \le B + c_{\mathsf{max}}$. Equivalently, we can show that $p(x^1) \le B + c_{\mathsf{max}}$

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.

and $p(x^2) \le B + c_{\mathsf{max}}$. 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 $p(x^1) \le B$. So suppose $x^1 \ne x^2$. 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. $x_j^1 \ne x_j^2$. Now, we prove the lemma by showing that

$$p(x_1) \le p(z) + \tilde{p}_j \cdot \frac{1-\tilde{x}_j}{\tilde{x}_j} \le B + c_{\mathsf{max}}. \tag{33}$$

Verify that the first inequality in (33) holds since x_1 and z are only different in their j-th entry: if $x_j^1 = 0$, then $p(x_1) \le p(z)$; 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, (33) is proved if we show that

$$p(z) \le B, \tag{34}$$

$$\tilde{p}_j(1-\tilde{x}_j)/\tilde{x}_j \le c_{\mathsf{max}}. \tag{35}$$

To verify (34), just note that $z \in \mathcal{P}$. To verify (35), 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:

$$\tilde{p}_j/\tilde{x}_j \le r u_j, \tag{36}$$

$$r \le \frac{c_j/u_j}{1-\tilde{x}_j}. \tag{37}$$

Observe that combining (36) and (37) implies (35). So, we are done if we prove (36) and (37) hold. To prove (36), it is enough to note that $Q_r(c_j/u_j) \le r\tilde{x}_j$, i.e. the area under the curve f_r that represents $Q_r(c_j/u_j)$ fits in a rectangle with width r and height $\tilde{x}_j$; this implies $\tilde{p}_j/\tilde{x}_j \le r u_j$. It is also straight-forward to verify (37) holds due to the concavity of f_r. □

G.3. Strictly Budget Feasible Mechanisms

Theorem 6 ensures that the budget constraint is not violated by an additive factor more than c_{max}. As we mentioned earlier, it is possible to convert the mechanism to a strictly budget feasible mechanism. First, we discuss a simple approach for guaranteeing strict budget feasibility that bears an arbitrary small loss in the competitive ratio. Then we discuss a more efficient approach that guarantees budget feasibility with no loss in the competitive ratio.

The approach with arbitrary small loss. When there is a known upper bound on c_{max}, namely $\overline{c_{\mathsf{max}}}$, we can simply run the mechanism with a reduced budget $B - \overline{c_{\mathsf{max}}}$. The competitive ratio remains the same (asymptotically) by the Small Bidders assumption. Even if $\overline{c_{\mathsf{max}}}$ is as large as ϵB for some $\epsilon > 0$, the competitive ratio would be at least $(1-1/e)(1-\epsilon)$. Note that $\overline{c_{\mathsf{max}}} = \epsilon B$ is always a valid upper bound for any constant $\epsilon > 0$, since having items with cost larger than ϵB is not consistent with the Small Bidders assumption.

The approach with no loss. We sketch the proof for the case when all utilities are equal to 1. The idea is not allocating from a single seller whose cost is equal to c_{max}, and reducing the budget to $B' = B - c_{\mathsf{max}}$. Then, by Property 3 of our rounding procedure, the payments would sum up to at most $B' + c_{\mathsf{max}} = B$. The mechanism will bear a loss of at most 1 in the total utility that it attains, which will not affect the competitive ratio by the Small Bidders assumption. The implementation details are discussed below.

When running Mechanism Truthful(f), after we finish running the auxiliary instance for a seller i and find the stopping rate r_i, we take an additional step to adjust the allocation rule offered to seller i. Let $c_{\mathsf{max}i} = \max\{c_j : j \neq i, j \in S\}$. We adjust the allocation rule so that it allocates 0 units from all costs above $c_{\mathsf{max}i}$. Note that the adjusted allocation rule is still monotone, and the (Myerson) payment to no seller will increase under this adjustment. With a proper tie-breaking rule (to break the possible ties between the sellers), the argument for the case of unit utilities is complete. The case of general utilities is similar: we will exclude the seller with the highest cost-utility rate.

Appendix H: Further discussion of the optimal allocation rule

We first provide the argument that shows proportional share mechanisms are not optimal. After that, in Sections H.2 and H.3, we study a stylized family of Bayesian instances with adversarial agents, and derive the optimal allocation rule for this family. We will see that the optimal allocation rule for this family is indeed equal to f^*, the optimal allocation rule in the prior-free setting. Our analysis will provide more intuition on our analysis of the prior-free setting. For example, in Section H.2 we will see how $\widehat{f}$ could be derived from the Lagrangian relaxation of the mathematical program that we write for the stylized Bayesian family.

In Section H.3, we provide more intuition on our impossibility instance from Section 7. We rediscover the distribution of costs in our impossibility instance (given by (2)) from the mathematical program written for the (adversarial) seller of our stylized Bayesian model. We provide an intuitive explanation for the solution of this program by connecting it to the the notion of *virtual costs*, a notion that might be of independent interest to readers familiar with the Myersonian approach in revenue maximizing auctions.

Finally, we emphasize that the goal in this section is merely providing more intuition on some of the tools that we used in the analysis of the prior-free setting. Our analysis here is done on a stylized class of Bayesian instances, and we do not expect that it could be used for providing alternative proofs of our main theorems.

H.1. Proportional Share Mechanisms

Define $f(d)=1$ for $d\leq 1$, and $f(d)=0$ for $d>1$. Here, we provide a proof sketch to show that Envy-Free(f) has competitive ratio 1/2 under the Small Bidders assumption. We omit the proof for truthfulness and budget feasibility.

Suppose that the set of sellers $S=\{1,\ldots,2n\}$, the budget $B=n$, and $u_i=1$ for all $i\in S$. Moreover, suppose that $c_i=0$ for all $i\leq n$ and $c_i=1$ for all $i>n$. We will verify that Envy-Free($\bar{f}$) has competitive ratio 1/2. First, see that in our instance, the scaling ratio of Envy-Free($\bar{f}$) is $r=1$: At $r=1$, the payment to seller i is 1 if $i\leq n$ and is 0 otherwise; these payments sum up to n (they exhaust the budget). Therefore, Envy-Free($\bar{f}$) buys the items of sellers $1,\ldots,n$ and attains utility n. On the other hand, note that the utility of the omniscient mechanism is $2n$, since the budget is equal to sum of the costs. The competitive ratio is therefore at most 1/2.

We also give a proof sketch to show the competitive ratio is at least 1/2 (without going into the details for tiebreaking). Suppose that Envy-Free($\bar{f}$) chooses a scaling ratio r, for a given instance. Let $T\subseteq S$ denote the set of items allocated by this mechanism, and let $U=\sum_{i\in T}u_i$. We will show that U^* is at most $2U$.

Under the Small Bidders assumption, we should have $\sum_{i\in T}ru_i\approx B$, since the mechanism roughly exhausts the budget. (The equality holds asymptotically, under the Small Bidders assumption. More precisely, we will have $B-c_{\mathsf{max}}\leq\sum_{i\in T}ru_i\leq B$.) On the other hand, for every $i\in S\backslash T$, we should have $\frac{c_i}{u_i}\geq r$. Therefore, the omniscient mechanism can attain utility at most B/r from the sellers in $S\backslash T$. So, $U^*\leq U+B/r$. But note that we showed $\sum_{i\in T}ru_i\approx B$. Consequently, $U^*\leq U+U$.

H.2. Rediscovering the optimal allocation rule in a stylized Bayesian setting

In this section, we consider a stylized Bayesian model with one buyer and one seller. The seller owns a divisible item with utility 1 and a cost which is drawn from a distribution with CDF C. The buyer's budget is equal to the mean of C. The buyer's goal is designing a mechanism which maximizes her utility subject to ex-post incentive compatibility, individually rationality, and ex-ante budget feasibility.[1] We suppose that after the buyer chooses an allocation rule f, the seller adversarially chooses a distribution C, with the goal of minimizing the buyer's utility.[2]

For a fixed f, we use P_f to denote the payment rule corresponding to f. For a fixed f, the seller's problem can be written as

$$\begin{aligned} &\min_C \int_\kappa f(\kappa)\,\mathrm{d}\,C(\kappa) \\ &\text{s.t} \int_\kappa P_f(\kappa)\,\mathrm{d}\,C(\kappa) = \int_\kappa \kappa\,\mathrm{d}\,C(\kappa). \end{aligned}$$

where κ is a parameter indexing cost. We then can write the Lagrangian as

$$\mathcal{L}(f, C, \lambda) = \int_\kappa [f(\kappa) - \lambda(P_f(\kappa) - \kappa)]\,\mathrm{d}\,C(\kappa). \tag{38}$$

The seller's and buyer's problems can then be written as $\max_\lambda \min_C \mathcal{L}(f, C, \lambda)$ and

$$\max_f \max_\lambda \min_C \mathcal{L}(f, C, \lambda), \tag{39}$$

respectively. The (unique) solution to (39) is characterized by the tuple (f^*, λ^*, C^*) where

$$f^*(\kappa) = \ln(e - \kappa), \quad \forall \kappa : 0 \le \kappa \le e - 1, \tag{40}$$

$$C^*(\kappa) = \frac{1}{e(1-\kappa)}, \quad \forall \kappa : 0 \le \kappa \le 1 - 1/e, \tag{41}$$

$$\lambda^* = 1/e.$$

This solution is unique up to the scaling of f. Note that this solution coincides with our results in Theorems 1 and 3. In fact, this result can also be proved as a corollary of Theorems 1, 2, and 3. We do not give a direct formal proof here. However, it is insightful to have a closer look at (39). We will find an intuitive connection between the optimality conditions from (39) and the fact that $\widehat{f^*}$ is a straight line. (We observed the latter fact in Proposition 2.)

[1] Our argument can be extended to a similar Bayesian instance with "many" iid sellers who own (possibly) indivisible items with unit utility.

[2] The buyer's budget is set to be equal to the mean of C. Alternatively, we can normalize the mean of C to 1.

First, let us write (39) in a slightly different form:

$$\max_{f,\lambda}\min_{C}\int_{\kappa}[f(\kappa)-\lambda(P_f(\kappa)-\kappa)]\,\mathrm{d}\,C(\kappa).$$

This form has a slightly different interpretation: the buyer chooses f, λ and then the seller responds to this choice with choosing C with the goal of minimizing $\mathcal{L}(f, C, \lambda)$. We investigate the buyer's choices for f, λ next.

Roughly speaking, for a "generic" choice of f, λ picked by the buyer, the seller has a unique best response: assuming that the expression $f(\kappa) - \lambda(P_f(\kappa) - \kappa)$ has a unique maximizer, namely at $\kappa = \kappa'$, the unique best response for the seller would be the degenerate distribution on the point κ'. The buyer, however, can choose f, λ in such a way that the seller becomes "indifferent" over all her choices: this can be done by choosing f, λ so that $f(\kappa) - \lambda(P_f(\kappa) - \kappa)$ is a constant for all possible κ, i.e.

$$f(\kappa)-\lambda(P_f(\kappa)-\kappa)\equiv\text{const},\quad\forall\kappa.$$

Such f can be characterized as the solution to the following differential equation:

$$\partial\frac{P_f(\kappa)-\kappa}{\partial f(\kappa)}=k_0,\tag{42}$$

where k_0 is a constant (with respect to κ). The solution to the above first-order linear ODE is given by

$$f(\kappa)=\log(k_0-\kappa/k_1)+k_2,$$

where k_1, k_2 are arbitrary constants. A careful investigation of the initial conditions gives us the solution that we expect: $f^*(\kappa) = \ln(e - \kappa)$ and $\lambda^* = 1/e$. As we mentioned before, this solution is unique up to the scaling of f^*. The scale parameter is in fact λ: for any $r > 0$, it is straight-forward to verify the tuple $(f^*_r, r\lambda^*, C^*_r)$ is indeed a solution to (39).[3]

Note that, by the definition of $\widehat{f}$, imposing (42) is just the same as imposing that $\widehat{f}$ is a straight line.

[3] Recall that f^*_r is just f^* scaled with ratio r. Consistent with the same notation, C^*_r denotes the CDF of the distribution that is obtained from scaling (stretching) the CDF C^* with ratio r along the horizontal axis.

In the above intuitive argument, we informally supposed that the optimality condition on the solution of (39), $f = f^*$, is given by choosing f^* in a way that the seller "becomes indifferent over her choices." We can formalize this argument and provide a direct way for solving (39). (The indirect way just uses our main theorems.) We omit the details.[4]

H.3. Rediscovering the impossibility instance in a stylized Bayesian setting

In this section, we provide further intuition on the impossibility instance of Section 7. Our approach is similar to the previous section, except that in here, we suppose that first the seller chooses C and then the buyer responds with choosing f. We will rederive the distribution of costs in the impossibility instance of Section 7 (given by (2)) from the mathematical program written for the (adversarial) seller.

Formally, consider an instance with one seller and one buyer. The seller owns a divisible item with utility 1 and a cost which is drawn from a distribution with CDF C. The buyer's budget, B, is supposed to be equal to the mean of C. After the seller chooses C (adverserially), the buyer responds with choosing a monotone allocation rule f. The buyer's goal is choosing an allocation rule which maximizes her total utility gain, subject to ex-post incentive compatibility, individually rationality, and ex-ante budget feasibility. The seller's goal is choosing C so that the buyer's total utility gain is minimized.

For a fixed C, the buyer's problem can be written as

$$\begin{aligned} \max \int_{\kappa} f(\kappa)\,\mathrm{d}\,C(\kappa) \\ \text{s.t} \int_{\kappa} P_f(\kappa)\,\mathrm{d}\,C(\kappa) = B, \\ f \text{ is monotone.} \end{aligned}$$

We can prove that the solution to this problem is indeed given by (40) and (41). In fact, this result can also be seen as a corollary of Theorems 1, 2, and 3. Rather than providing a full direct proof,

[4] The details would be made available to the interested reader upon request.

here we give a proof sketch that provides further insights by making a connection to the notion of *virtual costs*.

Using the fact that

$$P_f(\kappa) = \kappa f(\kappa) + \int_\kappa^\infty f(x)\,\mathrm{d}\,x,$$

we can rewrite the buyer's problem as

$$\begin{aligned} \max \int_\kappa & f(\kappa)\,\mathrm{d}\,C(\kappa) \\ \text{s.t} \int_\kappa & \left(\kappa f(\kappa) + \int_\kappa^\infty f(x)\,\mathrm{d}\,x\right)\mathrm{d}\,C(\kappa) = B, \\ & f \text{ is monotone.} \end{aligned}$$

Integration by parts allows us to rewrite the buyer's problem as

$$\begin{aligned} \max \int_\kappa & f(\kappa)C'(\kappa)\,\mathrm{d}\,\kappa \\ \text{s.t} \int_\kappa & f(\kappa)C'(\kappa)\cdot\left(\kappa + \frac{C(\kappa)}{C'(\kappa)}\right)\mathrm{d}\,\kappa = B, \\ & f \text{ is monotone.} \end{aligned}$$

The above mathematical program can be interpreted as the well-known knapsack problem as follows. There is a mass of items indexed by variable κ, which are to be packed in a knapsack with capacity B. Item κ has "volume" $\kappa + \frac{C(\kappa)}{C'(\kappa)}$ and gives a per-unit "felicity" 1. The density of items of type κ is $C'(\kappa)$.[5]

The usual approach to the knapsack problem is sorting items in an increasing order with respect to the volume-felicity ratio, and picking items in that order until the knapsack is full. Here, the volume-felicity ratio for item κ is equal to $\kappa + \frac{C(\kappa)}{C'(\kappa)}$, which we denote by $\phi(\kappa)$ and call it the *virtual cost* of item κ. One way to interpret the virtual cost term is that it encodes the cost of incentive constraints by adding a term $\frac{C(\kappa)}{C'(\kappa)}$ to the cost κ. This term resonates with the notion of "information rent" in forward auctions.

[5] Roughly speaking, one can think of the density of items of type κ as the "number" of items available of that type.

In the rest of this section, we use this new interpretation of the buyer's problem to provide an intuitive argument for showing that the optimal solution to the buyer's problem is indeed given by (41).[6] This argument is potentially interesting to readers familiar with the Myersonian approach in revenue maximizing forward auctions.

As we mentioned above, the optimal solution to the knapsack problem is attained by sorting the items with respect to the virtual cost term in an increasing order and picking the items in that order until the knapsack is full. The buyer can use this solution so long as monotonicity of f is not violated. We dismiss the monotonicity condition for the moment and offer an intuitive way for finding C^*. Then, we will see that the monotonicity condition is not binding at C^*.

For a "generic" choice of C, the knapsack problem has a generically unique solution, which is obtained by ordering items with respect to their virtual cost: items with lower virtual cost are more desirable to the buyer. An intuitive way to think about the solution of the minimax problem is that the (adversarial) seller chooses C in a way that all items become equally undesirable to the buyer: this happens when the virtual costs of all items are equal. That is, there exists a constant μ such that

$$\kappa + \frac{C(\kappa)}{C'(\kappa)} \equiv \mu. \tag{43}$$

Note that this condition is identical to the condition that our impossibility instance satisfies in Section 7. Recall from Figure 5 that the condition there is given by $\partial \frac{\kappa C(\kappa)}{\partial C(\kappa)} \equiv \mu'$, where the constant μ' represents the slope of the line in that figure. We can verify that this condition is equivalent to condition (43):

$$\begin{aligned} &\partial \frac{\kappa C(\kappa)}{\partial C(\kappa)} \equiv \mu' \\ \Leftrightarrow\ &\frac{\kappa C'(\kappa) + C(\kappa)}{C'(\kappa)} \equiv \mu' \\ \Leftrightarrow\ &\kappa + \frac{C(\kappa)}{C'(\kappa)} \equiv \mu'. \end{aligned}$$

[6] As we mentioned earlier, not only we can formalize this argument as a direct proof, but we can also provide an (almost immediate) indirect proof using Theorems 1, 2, and 3.

In the above argument, we informally supposed that the optimality condition is given by "making the buyer indifferent over all the items." With a little extra work, we can formalize this argument and provide a direct proof for solving the mathematical program and finding C^* (rather than doing it through Theorems 1, 2, and 3). Perhaps not surprisingly, this direct proof uses ideas similar to Section 7; in particular, it will use the two-dimensional plane that we saw in Figure 5. We omit the technical details, as our goal in this section is merely providing more intuition.[7] Finally, we remark again that the argument provided in this section works only for the discussed Bayesian setting, and we do not expect that it could be used for providing alternative proofs of our main theorems.[8]

Appendix I: 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 $F: 2^S \to \mathbb{R}^+$ defines the utility that the buyer derives from buying a subset of S. The buyer's 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-competitive mechanism is presented for it. Later, Chen et al. (2011) improved this result by giving an exponential-time deterministic mechanism with competitive ratio 0.119 and a polynomial-time randomized mechanism with competitive ratio 0.126.

Our Results. We study the problem under the Small Bidders assumption (Roughly speaking, that is when each individual can not significantly affect the market. See Section I.1 for a formal definition). Under this assumption, we design a deterministic mechanism which has competitive ratio $\frac{1}{2}$, but may have 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

[7] The details would be made available to the interested reader upon request.

[8] Recall that our theorems are for a prior-free setting without any distributional assumptions on costs or utilities.

functions. In fact, our mechanism is also a polynomial-time $(\gamma^2/2)$-competitive mechanism when it has access to a γ-approximation oracle for solving the knapsack problem (see Section I.2). To the extent of our knowledge, the best existing approximation oracle has $\gamma = 1 - 1/e$ due to Sviridenko (2004); this provides us a polynomial-time mechanism with competitive ratio $\gamma^2/2 \approx 0.2$.

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

Oneway-Truthfulness. All of the mechanisms that we design 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 I.1.3.

In Section I.4, we convert our oneway-truthful mechanisms to (fully) truthful mechanisms only by changing the payment rule. It is worth pointing out that computing the competitive ratio for oneway-truthful mechanisms is not more difficult than doing so for truthful mechanisms: Since the cost of optimum solution may only decrease if players report lower costs, then any α-competitive solution for the reported instance (possibly with costs reported lower than the actual costs) is also an α-competitive solution for the actual instance.

I.1. Preliminaries

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

I.1.1. Basic Definitions 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^*.

I.1.2. The Small Bidders Assumption Our Small Bidders assumption here is almost identical to the alternative Small Bidders assumption that was discussed in Section 2.1. Intuitively, it says that no individual can affect the (optimum solution of the) market significantly. This assumption is formally defined below.

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

Definition 8. We say that a market is *large* if $u_{\mathsf{max}} \ll U^*$.

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

I.1.3. 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 9. 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.

I.1.4. Submodular Functions Given the submodular function $F : 2^S \to \mathbb{R}^+$, we define an ordering of the elements of S with respect to F, which we call the *greedy sequence* and denote it by $\chi(F) = \langle x_1, \ldots, x_n \rangle$. For simplicity in the definition, we first define an auxiliary notion as follows: let $\chi_i = \cup_{j=1}^{i} \{x_i\}$ for all i, and let $\chi_0 = \emptyset$. The sequence is constructed such that

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

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 $x_1, \ldots, x_n$ 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 $i \leq n$.

I.2. The Exp-Time Mechanism

In this Section, we present a 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-utility rate*. We also call r^* the *optimum rate* when there is no risk of confusion.

The Selection Rule. 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 (i.e. the set of sellers whose items will be bought).

The Payment Rule. For simplicity, assume that the winners are indexed from $1, \ldots, k$. Payment to winner i is equal to $2r_i \cdot \partial_i$, where r_i is the optimum cost-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-utility rate in this instance.

Below we prove that the Oracle Mechanism is individually rational, oneway-truthful, and it has competitive ratio $\frac{1}{2}$. 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 competitive ratio of the mechanism asymptotically due to the Small Bidders assumption. The formal proof for this is deferred to Section I.5.

Mechanism 1: Oracle Mechanism

input : Submodular utility function F, Budget B

Sellers report their costs;

Compute the optimum utility F^* and the optimum cost-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**

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

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

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

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

end

Simplifying Assumption. Through out the analysis, w.l.o.g. we assume that $x_i = i$ for all $i \in S$.

LEMMA 22. *The Oracle Mechanism is individually rational.*

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

$$F(S^* \cup \chi_k) - F(\chi_k) \geq \frac{F^*}{2}. \tag{44}$$

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 (44) together imply that

$$c(S^* \cup \chi_k) - c(\chi_k) \geq \frac{c_k}{\partial_k} \cdot \frac{F^*}{2}. \tag{45}$$

On the other hand, recall that $c_k > 2r^* \cdot \partial_k$, so we can write (45) 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^*. □

LEMMA 23. *The Oracle Mechanism is oneway-truthful.*

Proof. For contradiction, suppose there exists a seller $i \in S$ who has incentive to report a cost $\overline{c_i}$ which is higher that 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 $\overline{c_i}$. 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. □

Lemma 24. *In a θ-large market, the Oracle Mechanism has competitive ratio $\frac{1}{2} - \theta$, i.e. asymptotically equal to $\frac{1}{2}$.*

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 Small Bidders assumption. □

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

As we mentioned before, under the Small Bidders assumption, we can convert any almost budget feasible mechanism to a budget feasible mechanism without any loss in its competitive ratio (asymptotically). This can be done simply by running the mechanism with a slightly reduced budget; the proof is deferred to Section I.5.

Lemma 25. *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 Small Bidders 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 $r_i \leq \frac{r^*}{1-\theta}$. □

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-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 11. Suppose we are given an instance of the problem with a submodular function F, budget B, and (publicly known) costs and utilities. A γ-*approximation oracle* is a polynomial-time algorithm that, for any instance of the problem, 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 7. *Suppose that Oracle Mechanism has access to a polynomial-time γ-approximation oracle for the knapsack optimization problem. Conditioned on using the γ-approximation oracle for computing the optimum cost-utility rates, the Oracle Mechanism becomes a polynomial-time $(\gamma/2)$-competitive 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 25. The key point in the proof of Lemma 25 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 25. 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 Small Bidders 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 $r_i \leq r^* \cdot (1-\theta)^{-1}/\gamma$. □

Corollary 3 **(of Theorem 7)**. *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 competitive ratio $\gamma^2/2$.*

I.3. An Improved Polynomial-Time Mechanism

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

Definition 12. 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 $\mathsf{r}(c, F, B)$, is then defined to be $B/F(\chi_k)$. We sometimes denote the stopping rate simply by $\mathsf{r}(B)$ when c, F are clear from the context.

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

The Selection Rule. 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 $r_i = \mathsf{r}(c', F, B/2)$ 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 defined to be 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 competitive 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 competitive ratio; the details are discussed in Sections I.4 and I.5.

Mechanism 2: A polynomial-time (1/3)-competitive 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**

  $c' \leftarrow c$;

  $c'_i \leftarrow \infty$;

  $r_i \leftarrow \mathsf{r}(c', F, B/2)$;

  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 with increasing index, i.e. $x_i = i$ for all $i \in S$.

LEMMA 26. *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 $\overline{\chi}(F) = \langle \overline{x}_1, \ldots, \overline{x}_{n-1} \rangle$ denote the greedy sequence for when i is removed. Also, let $\overline{\chi}_j$ denote the subset containing first j elements of $\overline{\chi}(F)$, and $\overline{\partial}_j = F(\overline{\chi}_j) - F(\overline{\chi}_{j-1})$.

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

First, suppose $\overline{k} = i-1$. 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 $\overline{k} \geq i$. Note that due to the definition of the greedy sequence we have $c_{x_i}/\partial_i \leq c_{\overline{x}_i}/\overline{\partial}_i$. So, the proof is done if we show that $r_i \geq c_{\overline{x}_i}/\overline{\partial}_i$. To this end, first observe that we have $F(\overline{\chi}_i) \cdot c_{\overline{x}_i}/\overline{\partial_i} \leq B/2$ due to the fact that $\overline{k} \geq i$. This just implies $c_{\overline{x}_i}/\overline{\partial_i} \leq B/(2F(\overline{\chi}_{\overline{k}}))$ due to the monotonicity of F. Finally, seeing that $B/(2F(\overline{\chi}_{\overline{k}})) = r_i$ proves the claim. □

LEMMA 27. *Mechanism 2 is oneway-truthful.*

Proof. The proof is almost identical to the proof of Lemma 23, we state the proof for completeness. For contradiction, suppose there exists a seller $i \in S$ who has incentive to report a cost $\overline{c_i}$ 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 $\overline{c_i}$. 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. □

Lemma 28. *Mechanism 2 has competitive ratio* $1/3$ *under the Small Bidders assumption.*

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 Small Bidders 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

$$F(S^* \cup \chi_{k+1}) - F(\chi_{k+1}) > \frac{2F^*}{3}. \tag{46}$$

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 (46) together imply that

$$c(S^* \cup \chi_{k+1}) - c(\chi_{k+1}) > \frac{c_{k+1}}{\partial_{k+1}} \cdot \frac{2F^*}{3}. \tag{47}$$

On the other hand, recall that $F(\chi_{k+1}) < F^*/3$, so we can write (47) 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^*. □

We prove that Mechanism 2 is almost budget feasible in Lemma 30 . Before that, we first need to prove the following lemma which will be used in the proof of Lemma 30. This Lemma states that, $F(\chi_k)$ is (roughly) at least half of the optimal utility.

LEMMA 29. *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 $\chi = \langle 1, \ldots, n \rangle$. 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 Small Bidders 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

$$F(S^* \cup \chi_{k+1}) - F(\chi_{k+1}) > \frac{F^*}{2}. \tag{48}$$

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 (48) 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^*.

□

Lemma 30. *Mechanism 2 is almost budget feasible*

Proof. Let r denote $\mathsf{r}(c,F,B/2)$, 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:

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

This would finish the proof due to the Small Bidders assumption.

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

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

$$\begin{aligned}
r_i = \frac{B/2}{F(\overline{\chi}_{\overline{k}})} &\leq \frac{B/2}{(1/2-\theta) \cdot F^*_i} && (49)\\
&\leq \frac{B/2}{(1/2-\theta) \cdot (1-\theta) \cdot F^*_{B/2}} && (50)\\
&\leq \frac{B/2}{(1/2-\theta) \cdot (1-\theta) \cdot F(\chi_k)}\\
&\leq \frac{r}{(1/2-\theta) \cdot (1-\theta)}\\
&\leq \frac{2r}{1-3\theta}
\end{aligned}$$

where (49) is due to Lemma 29 and (50) is due to the θ-largeness of the market. □

I.4. 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 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 competitive ratio remain the same since we have not changed the 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 31. *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 $\overline{c_i} > 2r_i \cdot \partial_i$, then it will not be selected as a winner anymore. The proof is by contradiction, suppose i reports a cost higher than $\overline{c_i} > 2r_i \cdot \partial_i$ and still gets selected as a winner.

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

Also, let $\overline{\chi}(F)$ denote the greedy sequence for the fake instance, $\overline{\chi}_j$ denote the subset containing first j elements of $\overline{\chi}(F)$, and $\overline{\partial}_j = F(\overline{\chi}_j) - F(\overline{\chi}_{j-1})$. Finally, let $\overline{S^*}$ 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 $F(\overline{S^*}) - F(\overline{\chi}_{\overline{k}}) \geq \frac{F_i^*}{2}$. Also, we will show that even if we get $\overline{\chi}_{\overline{k}}$ for free, the cost for buying an additional $F(\overline{S^*}) - F(\overline{\chi}_{\overline{k}})$ units of utility would be more than $2r_i$ *per unit of utility.* The two latter facts together would imply that

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

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

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

$$F(\overline{S^*} \cup \overline{\chi}_{\overline{k}}) - F(\overline{\chi}_{\overline{k}}) \geq \frac{F_i^*}{2}. \tag{51}$$

Now observe that by the definition of the greedy sequence $\overline{\chi}(F)$, the cost for buying each extra unit of utility, conditioned on having the subset $\overline{\chi}_{\overline{k}}$, is at least $\overline{c_i}/\overline{\partial}_{\overline{k}}$, i.e. the cost of the seller at position $\overline{k}$ of $\overline{\chi}(F)$ divided by the marginal utility that it adds. Also, see that $\overline{c_i}/\overline{\partial}_{\overline{k}} > 2r_i$ since $\overline{c_i} > 2r_i \cdot \partial_i$ and $\overline{\partial}_{\overline{k}} \leq \partial_i$. This proves that the cost for buying each extra unit of utility, conditioned on having the subset $\overline{\chi}_{\overline{k}}$, is more than $2r_i$. This fact, and (51) together imply that

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

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

Now, we prove the counterpart of Lemma 31 for Mechanism 2.

Lemma 32. *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 $\overline{c}$ denote the cost vector in the fake instance which is identical to the original cost vector c except that $\overline{c}_i$ is set to x. For any subset of sellers $S' \subseteq S$, let $\overline{c}(S') = \sum_{s \in S'} \overline{c}_s$; 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 $\overline{\overline{c}}$ denote the cost vector in the large instance which is identical to the original cost vector c except that $\overline{\overline{c}}_i$ is set to ∞. For any subset of sellers $S' \subseteq S$, let $\overline{\overline{c}}(S') = \sum_{s \in S'} \overline{\overline{c}}_s$.

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

Suppose that Mechanism 2 picks $\overline{\chi}_{\overline{k}}$ as the set of winners in the fake instance and $\overline{\overline{\chi}}_{\overline{\overline{k}}}$ as the set of winners in the large instance, i.e. the mechanism picks the first $\overline{k}$ elements of $\overline{\chi}(F)$ and the first $\overline{\overline{k}}$ elements of $\overline{\overline{\chi}}(F)$ 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 6. *The first $\overline{\overline{k}}$ elements of $\overline{\chi}(F)$ and $\overline{\overline{\chi}}(F)$ are identical.*

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

$$c(\overline{\overline{x}}_{\overline{\overline{k}}})/\overline{\overline{\partial}}_{\overline{\overline{k}}} \leq r_i, \tag{52}$$

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

On the other hand, see that

$$\begin{aligned} \overline{c}(x_{\overline{i}})/\overline{\partial}_{\overline{i}} &> r_i \cdot \partial_i/\overline{\partial}_{\overline{i}} \\ &\geq r_i \end{aligned} \tag{53}$$

where (53) is due to the fact that $\partial_i \geq \overline{\partial}_{\overline{i}}$; this fact holds because the first $i-1$ elements of $\chi(F)$ and $\overline{\chi}(F)$ are identical, which means the marginal utility added by seller i in $\chi(F)$ is more than her marginal in $\overline{\chi}(F)$.

To summarize, see that by (52) and (53) we have:

$$c(\overline{\overline{x}}_{\overline{\overline{k}}})/\overline{\overline{\partial}}_{\overline{\overline{k}}} \leq r_i,$$

$$\overline{c}(\overline{x}_{\overline{i}})/\overline{\partial}_{\overline{i}} > r_i$$

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

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

$$\begin{aligned} B/2 \geq & F(\overline{\chi}_{\overline{k}}) \cdot \overline{c}_{\overline{k}}/\overline{\partial_{\overline{k}}} && (54) \\ > & F(\overline{\chi}_{\overline{k}}) \cdot \frac{r_i \cdot \partial_i}{\overline{\partial}_{\overline{i}}} && (55) \\ \geq & F(\overline{\chi}_{\overline{k}}) \cdot r_i && (56) \\ \geq & F(\overline{\overline{\chi}}_{\overline{\overline{k}}}) \cdot r_i && (57) \\ = & B/2 \end{aligned}$$

(54) is due to the definition of $\overline{k}$ in the fake instance;

(55) holds since we have assumed $\overline{i} \leq \overline{k}$ (i.e. i wins in the fake instance) and so we have $\overline{c}(\overline{x}_{\overline{i}})/\overline{\partial_{\overline{i}}} \leq \overline{c}_{\overline{k}}/\overline{\partial_{\overline{k}}}$ by the definition of the greedy sequence, this fact, and the fact that $\overline{c}(\overline{x}_{\overline{i}}) > r_i \cdot \partial_i$ directly imply (55);

(56) holds since $\partial_i \geq \overline{\partial}_{\overline{i}}$; we already proved this is true in the proof of Claim 6: because the first $i-1$ elements of $\chi(F)$ and $\overline{\chi}(F)$ are identical;

(57) follows by the monotonicity of F and since $\overline{\overline{\chi}}_{\overline{\overline{k}}} \subseteq \overline{\chi}_{\overline{k}}$ (which holds by Claim 6). □

I.5. 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 25, 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 competitive ratio significantly. To this end, first we prove that the budget reduction does not affect the optimum utility significantly.

Lemma 33. *Assume we are given a θ-large market and Let F_b^* denote the optimal solution for this market when 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 $S^* = \{1, \ldots, s\}$. 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 $\chi(G) = \langle 1, \ldots, s \rangle$, were 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 $S' = \{1, \ldots, s'-1\}$. 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 Small Bidders assumption. □

To achieve a strictly budget feasible mechanism under the Small Bidders assumption, fix ϵ to be an arbitrary small constant. In a 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 25. Also, by Lemma 33, such a market (with the reduced budget) would be $\frac{\theta}{1-\theta-\epsilon}$-large. So, by Lemmas 24 and 33, the competitive ratio of the Oracle mechanism would be

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

See that when $\theta \to 0$, the competitive ratio would be

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

i.e. for any arbitrary small $\epsilon > 0$, we have a mechanism with competitive 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 competitive ratio $\frac{1}{3} \cdot (1-\epsilon)$.

Appendix J: Hoeffding Bounds

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

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

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

Endnotes

1. We note that the knapsack problem is solvable in polynomial time for the case of divisible items. It is not polynomial time solvable for indivisible items (unless P=NP), however, its optimum solution can be approximated with a multiplicative error $1+\epsilon$ in time polynomial in $n, \frac{1}{\epsilon}$.

2. Note that no mechanism can achieve utility larger than $U^*(\mathbf{c}, \mathbf{u})$.

3. We did not try to optimize the dependence on θ in our analysis.

4. In Corollary 2 in Section F, we compute a lower bound on the competitive ratio for the case of general utilities. This lower bound at , e.g., $\theta = 1/20$ and $\theta = 1/40$ is equal to 0.57 and 0.60, respectively.

5. This can be verified in an instance with a single seller.

6. A log-concave function is a function whose logarithm is concave. For example, any concave function is log-concave.

7. We can also prove Lemma 1 without computing this competitive ratio (see Appendix D). However, we choose this approach since it also determines the performance of other mechanisms in this family, which will also come in handy in proving the uniqueness result.

8. An explicit definition is $\widehat{f}(x) = Q_f(f^{-1}(x)) - f^{-1}(x)$.

9. This is so because an α-competitive mechanism in the prior-free setting is also α-competitive in Bayesian setting

10. Note that π is in fact a function of $\mathbf{c}_{-i}$. For simplicity, we safely suppress $\mathbf{c}_{-i}$ from the notation.

11. A proper tiebreaking rule is also needed for breaking the ties between items with the same cost-utility rate

12. The mechanism that we propose is discriminatory in the sense that it does not pay all bidders at the same rate. Such type of discrimination is natural in the applications that we discuss (see Section 10). For instance, in emission reduction auctions, firms which incur higher cost per unit of reduction have a higher reservation price, and will be paid at a higher rate per unit of reduction. (Note that although they will be paid at a higher rate, they sell necessarily a smaller number of units than firms with lower cost per unit of reduction.)

13. Because its total utility is equal to the total utility of the fractional solution.

14. David et al. (1999) contains a survey of arguments in favor and against public R&D funding.

15. To get some intuition on why this is a reasonable assumption , see, e.g., statistics of the auctions held by GreenInitiative (2015), and in particular, note the number of bidders, their bid prices, and bid quantities.

16. The validity of this assumption also depends on the individuals in the consideration set.

17. We 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 $\epsilon > 0$. The (asymptotic) competitive ratio will not be affected under the Small Bidders assumption.