

Learning on a Budget: Posted Price Mechanisms for Online Procurement

Ashwinkumar Badanidiyuru, Cornell University

Robert Kleinberg, Cornell University

Yaron Singer, Google

We study online procurement markets where agents arrive in a sequential order and a mechanism must make an irrevocable decision whether or not to procure the service as the agent arrives. Our mechanisms are subject to a budget constraint and are designed for stochastic settings in which the bidders are either identically distributed or, more generally, permuted in random order. Thus, the problems we study contribute to the literature on budget-feasible mechanisms as well as the literature on secretary problems and online learning in auctions.

Our main positive results are as follows. We present a constant-competitive posted price mechanism when agents are identically distributed and the buyer has a symmetric submodular utility function. For nonsymmetric submodular utilities, under the random ordering assumption we give a posted price mechanism that is $O(\log n)$ -competitive and a truthful mechanism that is $O(1)$ -competitive but uses bidding rather than posted pricing.

1. INTRODUCTION

In an online procurement market agents arrive in a sequential order to offer their service, and a buyer decides whether or not to procure the service as the agent arrives. Naturally, each agent wishes to be compensated for her service, while the buyer in turn has limited resources and aims to procure the services that maximize her utility.

The complexities of such markets originate in the severe informational limitations of the buyer: the agents' true costs for providing their service are unknown to the buyer and may be exaggerated, and the sequential manner in which agents appear requires making decisions without full knowledge of the market. In addition, the mechanism may be limited in the way it can collect information from agents, and must infer their costs from implicit signals.

To address the first informational limitation, mechanism design theory advocates for *truthful* mechanisms, which — through careful design of payment and allocation rules — incentivize every participating agent to reveal her true cost for performing the service. In procurement markets, the challenge is then to design truthful mechanisms that yield a favorable outcome for the buyer under limited resources. That is, optimize the buyer's utility under some fixed budget.

The framework for mechanism design under a budget is known as *Budget Feasibility* [Singer 2010] where the goal is to design truthful mechanisms that optimize the utility function of the buyer under the constraint that the total payments made to the agents do not exceed a fixed budget. Despite numerous impossibility results regard-

Ashwinkumar Badanidiyuru (ashwinkumarbv@gmail.com) and Robert Kleinberg (rdk@cs.cornell.edu), Computer Science Department, Cornell University; Yaron Singer (yarons@google.com), Google. Part of this work was conducted while Y.S. was visiting Cornell University.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2012 ACM 0000-0000/2012/02-ARTX \$10.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

ing overpayment of truthful mechanisms, the budget feasibility framework enables designing mechanisms with favorable *approximation* guarantees for broad classes of utility functions (see e.g. [Singer 2010; Chen et al. 2011; Dobzinski et al. 2011; Bei et al. 2011]).

The second informational limitation is due to the online arrival of the agents. Following the literature on online algorithms, we can make three main distinctions in modeling the way in which the arrival sequence of the agents is determined. The most general is the adversarial model, where an adversary chooses the arrival sequence of the agents and their cost in a manner that yields the worst possible outcome for the mechanism. Although in principal we would like our mechanisms to be robust to such input, such a strong assumption leaves little hope for making formal performance guarantees. A popular alternative is the *secretary model* which assumes the values are chosen adversarially, though the arrival order is chosen uniformly at random from the set of all possible permutations of the agents. This is the model considered by Dynkin [Dynkin 1963] for choosing the element with the highest value from an online sequence, and has been widely used since then in various other problem domains (see e.g. [Babaioff et al. 2007]). Lastly, a slightly stronger assumption is that the each agent is independently drawn from some unknown common distribution (see e.g. [Kleinberg and Leighton 2003; Babaioff et al. 2011; Besbes and Zeevi 2009]).

A procurement market can therefore be defined in terms of the objective function the buyer wishes to optimize, and the arrival model of the agents. The goal is to design truthful budget feasible mechanisms that maximize the buyer's utility. There are two common approaches for agents to express their cost. A *bidding* mechanism learns the costs by soliciting a bid from each agent upon its arrival. The bidding model is often considered unnecessarily complex, especially from the standpoint of the bidders, and a preferred alternative is the *posted price* mechanism. A posted price mechanism presents each agent with a (possibly different) price, and the agent can either accept or reject the mechanism's offer. Posted price mechanisms are compelling due to their simplicity and are the most commonly used form of pricing. As such, there is large body of algorithmic research on posted price mechanisms (see e.g. [Chawla et al. 2010]).

Naturally, the bidding model is stronger than the posted price model, since a mechanism that solicits bids can always simulate a posted price mechanism by deciding on an offer price without observing the agent's bid, and pay the agent only if her bid does not exceed the offer. There are cases, however, where bids cannot be solicited and only posted price mechanisms can be implemented. Will the market suffer? This is the main question we wish to address in this paper:

Are posted prices as powerful as bidding in online procurement?

1.1. Posted Price Mechanisms in Online Procurement

In this paper we study posted price mechanisms in online procurement markets. We address the question above by studying various online procurement markets characterized by the buyer's utility and the arrival order of the agents. It is easy to show that the most general class of objectives for which budget feasible mechanisms can be obtained is that of subadditive utility functions, and when taking computation limitations into account, the most general class known is the class of *nondecreasing submodular* utility functions. We will therefore focus our attention on this class.

The main result in this paper shows that there are procurement markets where posted price mechanisms are competitive with bidding mechanisms.

THEOREM. [INFORMAL:] *In any symmetric submodular procurement market where the agents arrive from an unknown distribution, posted price mechanisms perform asymptotically as well as bidding mechanisms.*

In particular, we show that there is a posted price universally truthful budget feasible mechanism which is $O(1)$ -competitive when agents are drawn i.i.d from an unknown distribution. This is quite surprising since posted price mechanisms are known to be severely limited in comparison to bidding mechanisms in auctions where agents' values are drawn from an unknown distribution [Babaioff et al. 2011].

For the general case of nondecreasing submodular markets with agents that arrive in the secretary model, we show that the performance of posted price mechanisms is at least a logarithmic (in the number of agents) approximation of bidding mechanisms: We present a posted price mechanism in this model which is $O(\log n)$ -competitive as well as a bidding mechanism which is $O(1)$ -competitive, both for nondecreasing submodular functions. These mechanisms are based on estimating appropriate thresholds which are used to decide which agents are allocated.

THEOREM. [INFORMAL:] *In any nondecreasing submodular procurement market where the agents' arrival order is chosen uniformly at random, posted price mechanisms can be at least a $O(\log n)$ approximation to bidding mechanisms.*

It is an open question whether there is indeed a logarithmic gap between the performance of bidding and posted price mechanisms in submodular procurement markets, since it may be possible that there are posted price mechanisms with better guarantees than the one shown here. We discuss this open question in our concluding remarks.

1.2. Organization of the Paper

We begin with several definitions and formalize the model in Section 2. In Section 3 we present a posted price mechanism for symmetric submodular procurement markets, where agents private values correspond to i.i.d draws from an unknown distribution. We then show a mechanism for the general case of nondecreasing submodular procurement markets in the secretary model in Section ???. We complement this result with a constant-competitive bidding mechanism in Section 5. We conclude with a brief discussion and highlight some open problems in Section 6.

2. DEFINITIONS

2.1. Model

We study procurement auctions where there is a set of n agents, denoted \mathcal{N} , each offering a service for which they associate a private cost $c_i \in \mathbb{R}_+$. The buyer has a public (known to the mechanism designer and maybe known to all agents) budget $B \in \mathbb{R}_+$ and a public nondecreasing utility function $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}_+$ over the subsets of agents.

An important element in our model is the online arrival of agents. We assume that there are n different time steps, in each step $i \in [n]$ a single agent appears and the mechanism makes a decision that is based on the information it has about the agent and the history of the previous $i - 1$ stages. In this paper we distinguish between three assumptions on how the order of agents is determined:

- (1) **The adversarial model:** The agents costs and their arrival order are chosen by an adversary. The adversary chooses the worst arrival order and costs and has full knowledge of the mechanism though it cannot observe the actions the mechanism takes (this is known as an *oblivious* adversary).

- (2) **The secretary model:** The agents costs are chosen by an adversary, though their arrival order is a permutation that is drawn uniformly at random from the set of all possible permutations over the agents.¹
- (3) **The i.i.d model:** At each time step, an agent is drawn from some unknown distribution. We use this model only for symmetric utility functions, i.e. functions for which the value only depends on the cardinality of the set. In such cases this model is equivalent to having a sequence of costs drawn from an unknown distribution.

The above models are described in a decreasing order of generality: the i.i.d. model is a special case of the secretary model which is a special case of the adversarial model.

2.2. Mechanisms

A mechanism $\mathcal{M} = (A, p)$ is a pair that consists of an allocation function² $A : \mathbb{R}_+^n \rightarrow 2^{[n]}$ and a payment function $p : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$. The allocation function selects a subset of agents $S = A(b)$ given a set of bids b , and we use $A_i(b)$ to denote the indicator function that returns 1 if agent a_i is allocated and 0 otherwise. The payment function returns a vector of payments that describes the agents' compensation, where $p_i(b)$ is the payment to agent a_i when the bid vector is b .

In the case where agents arrive in an online fashion, as in our model, the input is received in a sequential manner and the mechanism decides on the allocation and payment at every stage i as agent a_i appears and places her bid, and the mechanism returns $(A_i(b_1, \dots, b_i), p_i(b_1, \dots, b_i))$. The allocation is the set $S = \{a_i : A_i(b_1, \dots, b_i) = 1\}$ and the total payments are $\sum_{i=1}^n p_i(b_1, \dots, b_i)$. For brevity, when it will be clear from the context, we will use p_i to denote the payment for agent a_i .

In our model we seek *truthful (incentive compatible)* mechanisms where reporting the true costs is a dominant strategy for agents. Formally, a mechanism $\mathcal{M} = (A, p)$ is *truthful* if for every $a_i \in \mathcal{N}$ with cost c_i and bid b_i , and every set of bids b_{-i} of the $\mathcal{N} \setminus \{a_i\}$ agents (or $\{a_1, \dots, a_{i-1}\}$ agents in the online arrival case) we have:

$$p_i(c_i, b_{-i}) - c_i \cdot A_i(c_i, b_{-i}) \geq p_i(b_i, b_{-i}) - c_i \cdot A_i(b_i, b_{-i})$$

A randomization over truthful mechanisms is a *universally truthful* mechanism. An important property we require is that the mechanism's payments should not exceed the budget: $\sum_i p_i \leq B$. This property is called *budget feasibility*. In addition, as common in mechanism design, we seek normalized ($a_i \notin S$ implies $p_i = 0$), individually rational ($p_i \geq c_i$) mechanisms with no positive transfers ($p_i \geq 0$).

In this paper we will be particularly interested in *posted price* mechanisms. Such mechanisms offer each agent a price (may offer different agents different prices), and the agent can accept or reject the mechanism's offer. Since agents are assumed to be rational, an agent accepts the price if her cost is below the price offered by the mechanism and rejects it otherwise. To each agent that accepts the offer, the mechanism must pay at least the price it was offered. We will also discuss the bidding model where the mechanism solicits bids from the agents as they arrive and makes an irrevocable decision on whether the agent is allocated and how much she is paid.

The objective is to maximize the utility function under the budget, i.e. allocate to the subset of agents S that yields the highest value possible, under all the constraints discussed above. We will compare our mechanisms against the most demanding bench-

¹This assumption on the arrival is known as the *secretary model* due to Dynkin's celebrated "secretary problem", where this assumption is introduced.

²Although the mechanism *procures* services from agents we use the term *allocation function* to be consistent with traditional mechanism design terminology. Since the mechanism allocates budget resources to the selected agents, this term seems appropriate in this case as well.

mark possible. We will compare against the optimal solution that is obtainable without computational limitations, when all the agents true costs are known in advance. Our goal is procure a set that has a value that is as close to the optimal solution as possible. Formally, for $\alpha \geq 1$ we say that a mechanism is α -competitive if it selects a set S such that in expectation over the arrival order of the agents and the randomization of the mechanism, we have that $f(S^*) \leq \alpha f(S)$, where S^* is the optimal solution. The mechanisms we construct have the additional property of being computable in polynomial time. In cases where the utility function requires exponential data to be represented, we take the common “black-box” approach and assume that f is represented by an oracle. We will use the *value query* model where given a set S , the oracle returns $f(S)$ (see [Blumrosen and Nisan 2005] for more details).

3. POSTED PRICE MECHANISMS FOR UNKNOWN DISTRIBUTIONS

In this section we assume that the utility function f is a symmetric submodular function, and we design a $O(1)$ -competitive posted price mechanism for the i.i.d. model. Our focus throughout most of this section is on the special case $f(S) = |S|$. In this special case, our goal can be restated as follows: we must show how a buyer with a budget of B can purchase approximately as many services as possible using a sequence of take-it-or-leave-it offers to sellers with i.i.d. costs, without knowing the distribution of costs in advance. Later, in Subsection 3.2, we show how this result for the function $f(S) = |S|$ easily implies a constant-factor approximation for general symmetric submodular f . The case of non-symmetric submodular functions is deferred to Sections 4 and 5. We emphasize that in this section we target only a constant-factor approximation, and we do not optimize for this constant.

It is instructive to compare our setting to an online single-item posted-price auction with i.i.d. bidders. For this problem [Babaioff et al. 2011] shows a $\Omega(\frac{\log(h)}{\log(\log(h))})$ lower bound. Here h is the ratio between maximum possible value and minimum possible value. Our result illustrates that the hardness of posted pricing in the i.i.d. model with an unknown distribution does not extend to budgeted procurement with a symmetric submodular utility function. To achieve a constant-approximation, we need to develop a novel algorithm that is quite different from existing algorithms in the literature on online posted pricing. Specifically, rather than reducing to a multi-armed bandit problem and using a general-purpose bandit algorithm (which is inapplicable in our setting because of the budget constraint) we design an algorithm that is quite different from bandit algorithms: in effect, it emulates a stochastic search for the optimal price using a multiplicative-increase multiplicative-decrease rule, with a simple sample-and-average policy guiding the multiplicative updates. We show, via a somewhat intricate analysis, that the stochastic search inflicts only a constant-factor loss relative to a buyer with foreknowledge of the optimal price. In effect, this involves showing that the algorithm has such a strong bias toward the optimal price that (with constant probability) a constant fraction of all the prices it offers are at or near the optimal price.

3.1. Algorithm for the Special Case $f(S) = |S|$

We present our algorithm in pseudocode, followed by a plain English description.

A Posted Price Mechanism for unknown distributions

Initialize the constants $\delta = 1/10$, $r = 2$, $a = 4000$, $z = 300000$.

With probability $1/2$:

Consider states p_1, p_2, \dots, p_m where $p_1 = \frac{B}{n}$, $p_{i+1} = r \cdot p_i$ for $1 \leq i < m$, and $\frac{B}{rz} < p_m \leq \frac{B}{z}$.
Initialize state $i = m$.

TEST(i): Set price $p = p_i$.

1. Offer price p to the next $\frac{anp}{B}$ sellers, or until the sellers or budget runs out.
In case more than $a(1 + \delta)$ offers are accepted, then stop offering and move to step 2.
2. Let t be the number of sellers who accept.
 - i. If $a(1 - \delta) \leq t \leq a(1 + \delta)$ then go to TEST(i).
 - ii. If $t > a(1 + \delta)$ then update i to $i - 1$ and go to TEST($i - 1$).
 - iii. If $t < a(1 - \delta)$ then update i to $i + 1$ and go to TEST($i + 1$).

With probability $1/2$:

Allocate to the first agent with cost $c_i \leq B$

The algorithm uses states $1, 2, \dots, m$ with associated prices p_1, p_2, \dots, p_m that form a geometric progression with common ratio r . The progression starts at B/n , so that $p_1 = B/n$, $B/(z \cdot r) < p_m \leq B/z$, and $p_{i+1} = r \cdot p_i$ for $0 \leq i < m$. The main loop of the algorithm is a loop that, in state i , runs a subroutine TEST(i) whose output is an element of $\{i - 1, i, i + 1\}$. The output of TEST(i) becomes the new state.

Subroutine TEST(i) operates as follows. It offers price $p = p_i$ to anp/B bidders, for some constant a . If the number of accepted offers is ever greater than $a(1 + \delta)$ it quits the subroutine immediately and outputs state $i + 1$. If the number of successes is less than $a(1 - \delta)$, then it outputs state $i - 1$. Otherwise it outputs state i .

Let ℓ be the index such that $B/p_{\ell-1} \geq n \cdot s(p_{\ell-1})$ and $B/p_\ell < n \cdot s(p_\ell)$ where $s(p)$ is the probability that a seller sells while offering at price p . (Note that B/p is a decreasing function of p whereas $s(p)$ is non-decreasing, hence p_ℓ is undefined if and only if $B/p_m \geq n \cdot s(p_m)$). But in such a case the second half of our algorithm — which with probability $1/2$ sells to a single player — is a $2r^2z$ approximation according to Lemma 3.1.)

LEMMA 3.1. $\mathbb{E}(|OPT|) \leq 2r \cdot B/p_\ell$

PROOF. Consider any specific realization of the random variables representing the costs of each seller. Define

$$\begin{aligned} O_1 &= \{i | c_i \leq p_{\ell-1} \text{ and } i \in OPT\} \\ O_2 &= \{i | c_i > p_{\ell-1} \text{ and } i \in OPT\} \\ O_3 &= \{i | c_i \leq p_{\ell-1}\}. \end{aligned}$$

We obtain our upper bound on $\mathbb{E}(|OPT|)$ by comparison with the expected cardinalities of O_1, O_2, O_3 as follows.

$$\begin{aligned} \mathbb{E}(|OPT|) &= \mathbb{E}(|O_1| + |O_2|) \\ &\leq \mathbb{E}(|O_3|) + \mathbb{E}(|O_2|) \\ &= ns(p_{\ell-1}) + \mathbb{E}(|O_2|) \\ &\leq ns(p_{\ell-1}) + \frac{B}{p_{\ell-1}} \\ &\leq \frac{B}{p_{\ell-1}} + \frac{B}{p_{\ell-1}} \\ &= 2r \frac{B}{p_\ell} \end{aligned}$$

□

Define the set of *good states* to be $GS = \{p_i | i \leq \ell \text{ and } s(p_i) \geq \frac{B(1-2\delta)}{np_i}\}$. Note that $p_\ell \in GS \subseteq \{p_{\ell-1}, p_\ell\}$. In a good state p_i , the expected number of accepted offers during $\text{TEST}(i)$ is

$$\frac{anp_i s(p_i)}{B} \geq (1-2\delta)a.$$

This partially justifies the term “good states”, since the algorithm is designed to find a state in which the expected number of accepted offers equals a .

LEMMA 3.2. *If $p = p_i$ satisfies $p < p', \forall p' \in GS$ then $\text{TEST}(i)$ outputs $i+1$ with probability at least $1 - (1-2\delta)/(\delta^2 a)$. If $p > p', \forall p' \in GS$ then $\text{TEST}(i)$ outputs $i-1$ with probability at least $1 - (1-2\delta)/(\delta^2 a)$.*

PROOF. We will use the fact that the variance of a Bernoulli random variable is bounded above by its expectation; consequently the same property holds for sums of independent Bernoulli random variables. By an application of Chebyshev’s Inequality, then, if the expectation of a sum of Bernoulli random variables is y , then the probability that the sum differs from its expectation by more than w is at most y/w^2 .

If $s(p) < \frac{(1-2\delta)B}{np}$ then the expected number of successes in $\text{TEST}(i)$ is less than $\frac{anp}{B} \cdot \frac{(1-2\delta)B}{np} = (1-2\delta)a$. Let $y \in [0, (1-2\delta)a]$ be the expected number of successes. The probability that the number of successes is greater than or equal to $a(1-\delta)$ is at most $\frac{y}{(a(1-\delta)-y)^2} \leq \frac{(1-2\delta)a}{\delta^2 a^2} = \frac{1-2\delta}{\delta^2 a}$.

If $s(p) \geq s(p_i) > \frac{B}{np_i} \geq \frac{rB}{np} > \frac{B}{(1-2\delta)np}$ then the expected number of successes in $\text{TEST}(i)$ is greater than $\frac{anp}{B} \cdot \frac{B}{(1-2\delta)np} = (1-2\delta)^{-1}a$. The probability that the number of successes is less than or equal to $a(1+\delta)$ is at most $\frac{a/(1-2\delta)}{(\delta+2\delta^2)^2 a^2 / (1-2\delta)^2} \leq \frac{1-2\delta}{\delta^2 a}$. \square

Informally, Lemma 3.2 says that when the current price is far from the good states, the random walk is biased to drift in the direction of the good states. Analyzing this biased random walk is the crux of our proof. To simplify the analysis, we couple the algorithm’s random walk with a simpler random walk in which the bias to drift toward the good states is exactly equal to the lower bound asserted in Lemma 3.2. Denote this bias by $\beta = 1 - (1-2\delta)/(\delta^2 a)$. The algorithm’s Markov chain \mathcal{M}_1 is defined to have states $1, 2, \dots, m$, and the transition probability from state i to state $j \in \{i-1, i, i+1\}$ is equal to the probability that an execution of $\text{TEST}(i)$ finishes by calling $\text{TEST}(j)$. (Conditional on having a nonzero budget and number of sellers remaining at the end of $\text{TEST}(i)$.) We compare this Markov chain to another one, \mathcal{M}_2 , defined on the set of non-negative integers. State 0 is an absorbing state of \mathcal{M}_2 and in every other state $i > 0$, the transition probabilities to states $i-1$ and $i+1$ are β and $1-\beta$, respectively. The two Markov chains are depicted in Figures 1 and 2.

For $1 \leq i \leq m$, let $\Delta(i)$ denote the distance from state i to the set of good states, i.e.

$$\Delta(i) = \min_{j \in GS} |i - j|.$$

The following lemma, stated informally, says the Markov chain \mathcal{M}_2 starting from state $\Delta(i)$ provides an upper bound on the stochastic process describing the distance from GS when one starts \mathcal{M}_1 in state i and runs it until it reaches GS .

LEMMA 3.3. *Let i be any state of \mathcal{M}_1 . There is a coupling of \mathcal{M}_1 and \mathcal{M}_2 such that \mathcal{M}_1 starts in state $i_0 = i$, \mathcal{M}_2 starts in state $k_0 = \Delta(i)$, and at any time $t \geq 0$ if the first Markov chain’s state sequence i_0, i_1, \dots, i_{t-1} does not include any element of GS , then the current pair of states (i_t, k_t) satisfies $k_t \geq \Delta(i_t)$.*

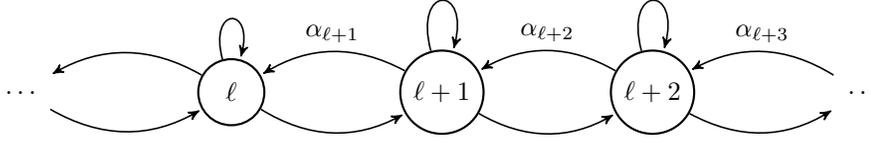
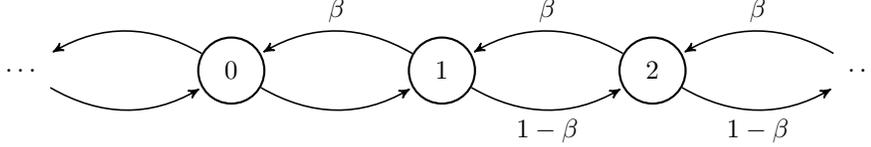
Fig. 1. Actual markov chain with $\alpha_i \geq \beta$ for $i \geq l+1$ 

Fig. 2. New markov chain

PROOF. The coupling is easy to describe. When the two Markov chains are in states (i, k) , respectively, if $i \in GS$ we couple them arbitrarily, e.g. by updating i and k independently using the transition probability of their respective Markov chains. If $i \notin GS$ let i' denote the neighboring state that is closer to GS than i : $i' = i + 1$ if $i < j$ for all $j \in GS$, and $i' = i - 1$ if $i > j$ for all $j \in GS$. Let α_i denote the transition probability from i to i' in \mathcal{M}_1 , and let σ_i denote the transition probability from i to itself in \mathcal{M}_1 . Our coupling in state (i, k) works as follows. With probability β we update i to i' and k to $k - 1$. With the remaining $1 - \beta$ probability, we update k to $k + 1$ and update i as follows: it transitions to i' with probability $\alpha_i - \beta$, to i with probability σ_i , and to $2i - i'$ (the other neighboring state) with probability $1 - \alpha_i - \sigma_i$. It is immediate from our definition of the coupling that if (i_{t-1}, k_{t-1}) and (i_t, k_t) are two consecutive state pairs such that $i_{t-1} \notin GS$, then

$$\Delta(i_t) - \Delta(i_{t-1}) \leq k_t - k_{t-1}. \quad (1)$$

The lemma follows by summing (1) over all time steps preceding t . \square

LEMMA 3.4. *Consider a sequence of phases in our algorithm's execution beginning with $TEST(i)$ where $i \notin GS$ and ending immediately before the first subsequent instance of $TEST(j)$ such that $j \in GS$. In total during this sequence of phases, the algorithm makes no more than $(anp_\ell/B)\phi(\Delta(i))$ offers in expectation and spends no more than $a(1 + \delta)p_\ell\phi(\Delta(i))$ of its budget in expectation, where $\phi(\cdot)$ is the function*

$$\phi(k) = \frac{r}{r - \beta - (1 - \beta)r^2} \cdot (r^k - 1). \quad (2)$$

PROOF. Let ℓ' denote the minimum element of GS (either ℓ or $\ell - 1$). For a given value Δ , there are at most two states i such that $\Delta(i) = \Delta$, namely $i = \ell + \Delta$ and $i = \ell' - \Delta$. In state $\ell + \Delta$, the algorithm makes $anp_i/B = anp_\ell r^\Delta/B$ offers in expectation, and it spends at most $a(1 + \delta)p_i = a(1 + \delta)p_\ell r^\Delta$ because subroutine $TEST(i)$ stops making offers after $a(1 + \delta)$ offers have been accepted. In state $\ell' - \Delta$ the algorithm makes $anp_\ell r^{-\Delta}/B$ offers in expectation, and it spends at most $a(1 + \delta)p_\ell r^{-\Delta}$. Thus, for a given value of Δ , the expected number of offers made and the expected amount spent in state $\ell + \Delta$ are both greater than their counterparts in state $\ell' - \Delta$. Accordingly, it suffices to prove the lemma for $i = \ell + \Delta$.

To do so, we use the coupling provided by Lemma 3.3. Our algorithm proceeds through a sequence of states i_0, i_1, \dots, i_T such that i_T is the first state in the sequence that belongs to GS . Meanwhile Markov chain \mathcal{M}_2 proceeds through a sequence

of states k_0, k_1, \dots, k_T such that $k_t \geq \Delta(i_t)$ for all $t = 0, \dots, T$. By the foregoing discussion, the expected number of offers made by our algorithm and the expected amount spent are respectively bounded above by

$$\begin{aligned} \mathbb{E}[\text{offers made}] &\leq \sum_{t=0}^{T-1} anp_\ell r^{k_t} / B = \frac{anp_\ell}{B} \sum_{t=0}^{T-1} r^{k_t} \\ \mathbb{E}[\text{amount spent}] &\leq \sum_{t=0}^{T-1} a(1+\delta)p_\ell r^{k_t} = a(1+\delta) \sum_{t=0}^{T-1} r^{k_t}. \end{aligned}$$

Thus, to complete the lemma, it suffices to bound $\mathbb{E} \left[\sum_{t=0}^{T-1} r^{k_t} \right]$. We do so by introducing another stopping time τ , defined to be the earliest t such that $k_t = 0$. Note that $\tau \geq T$, so $\sum_{t=0}^{\tau-1} r^{k_t} \geq \sum_{t=0}^{T-1} r^{k_t}$. Define

$$\phi(k) = \mathbb{E} \left[\sum_{t=0}^{\tau-1} r^{k_t} \mid k_0 = k \right].$$

The function ϕ satisfies the linear recurrence

$$\phi(k) = r^k + \beta\phi(k-1) + (1-\beta)\phi(k+1)$$

for $k > 0$, with the initial condition $\phi(0) = 0$. Solving the recurrence we find that equation (2) in the statement of the lemma specifies the solution. \square

In the sequel, let $c_1 = \frac{\tau}{r-\beta-(1-\beta)r^2}$.

COROLLARY 3.5. *Consider the sequence of phases beginning with the algorithm's initialization and ending with the first subsequent instance of TEST(j) such that $j \in GS$. In total during this sequence of phases, the algorithm makes no more than $(ac_1/z) \cdot n$ offers in expectation and spends no more than $(a(1+\delta)c_1/z) \cdot B$ of its budget in expectation,*

PROOF. The algorithm begins in phase m . Using the fact that $r^{m-1} \leq n/z$, this implies the bound $\phi(\Delta(m)) = c_1(r^{m-\ell} - 1) < (c_1 r^{1-\ell})n/z$. By Lemma 3.4, the expected number of offers made before reaching a good state is at most

$$\frac{anp_\ell}{B} (c_1 r^{1-\ell}) (n/z) = \frac{anr^{\ell-1}(B/n)}{B} (c_1 r^{1-\ell}) (n/z) = \left(\frac{ac_1}{z}\right) n.$$

The bound in the expected amount spent follows by an analogous calculation:

$$a(1+\delta)p_\ell(c_1 r^{1-\ell})(n/z) = a(1+\delta)r^{\ell-1}(B/n)c_1 r^{1-\ell}n/z = (a(1+\delta)c_1/z) \cdot B.$$

\square

THEOREM 3.6. *The algorithm is constant-competitive. In fact, with probability at least $\frac{1}{2}$, the number of offers accepted is at least $c \cdot (B/p_\ell)$, where c is an absolute constant.*

PROOF. We distinguish two cases. If $B/p_\ell \leq 7 \times 10^8$ then with probability 1/2, we expend our full budget on a single agent. When this happens, the number of offers accepted equals 1, which is at least $(1.4 \times 10^{-9})(B/p_\ell)$.

If $B/p_\ell \geq 7 \times 10^8$ then we will show that with constant probability, the algorithm executes enough calls to TEST(i) for states $i \in GS$ that it succeeds in getting $c \cdot (B/p_\ell)$ offers accepted. We will define five bad events E_1, \dots, E_5 , each having probability at most 0.1, such that the number of offers accepted is greater than $c \cdot (B/p_\ell)$ whenever none of the events E_1, \dots, E_5 occurs. Corollary 3.5 presents bounds on the expected number of offers made, and the expected amount spent, during the ‘‘startup stage’’

before the first time the algorithm reaches a good state. Let E_1 (resp. E_2) denote the event that the actual number of offers made (resp. amount spent) in the startup stage exceeds the bound given in Corollary 3.5 by a factor of more than 10. By Markov's inequality, each of E_1, E_2 has probability bounded by 0.1.

Imagine running the algorithm for an infinite number of steps, disregarding the fact that it eventually spends more than its budget and makes more than its allotted n offers. Define a *good phase* to be an execution of $\text{TEST}(i)$ such that $i \in GS$. Define an *excursion* to be the (possibly empty) set of offers made between two consecutive good phases. Let

$$x = \min \left\{ \frac{B}{4p_\ell(a + 10ac_1)}, \frac{B}{40(1 + \delta)ac_1p_\ell} \right\}$$

and consider the first x good phases along with the x excursions that occur immediately after each of them. How many offers, in expectation, are made during these x phases and excursions? The expected number of offers during the good phases is bounded above by $(anp_\ell/B) \cdot x$. According to Lemma 3.4 the expected number of offers during the excursions is bounded above by $(anp_\ell/B)c_1(r - 1)$ and the expected amount spent during the excursions is bounded above by $a(1 + \delta)p_\ell c_1(r - 1)$. Let E_3 (resp. E_4) denote the event that the actual number of offers (resp. amount spent) made during the excursions does not exceed this bound by a factor of more than 10. Once again, by Markov's inequality, each of E_3, E_4 has probability bounded by 0.1.

Our choice of the parameters a, r, z, δ has been designed to ensure that, assuming events E_1, E_2, E_3, E_4 do not occur, the combined number of offers made until the end of the x^{th} good phase is less than n and the combined amount spent until this time is less than B . Thus, all of the offers made during the first x good phases actually took place during the algorithm's execution, i.e. before the budget was expended or the maximum number of offers was reached. In any good phase, the expected number of offers accepted is at least $a(1 - 2\delta)$. Let E_5 denote the event that fewer than $a(1 - 3\delta)x$ offers are accepted in all of the first x good phases combined. Applying Chebyshev's Inequality to this sum of independent Bernoulli random variables as in the proof of Lemma 3.2 implies that $\Pr(E_5) \leq 0.1$, again by our choice of a, r, z, δ .

By the union bound, the probability that none of E_1, \dots, E_5 occur is at least $\frac{1}{2}$, and when this happens at least $a(1 - 3\delta)x$ offers are accepted. The theorem follows, because $x = \Omega(B/p_\ell)$. \square

3.2. Extension to Symmetric Submodular Functions

When f is a symmetric submodular function, it means that there exists a nondecreasing concave function g such that $f(S) = g(|S|)$ for all sets S . It turns out that the argument from Section 3.1 carries through to this case with very few modifications. As an upper bound on $\mathbb{E}[f(OPT)]$, we use the following lemma which generalizes Lemma 3.1.

LEMMA 3.7. $\mathbb{E}[f(OPT)] \leq 2r \cdot g(B/p_\ell)$

PROOF. Consider any specific realization of the random variables representing the costs of each seller. Define

$$\begin{aligned} O_1 &= \{i | c_i \leq p_{\ell-1} \text{ and } i \in OPT\} \\ O_2 &= \{i | c_i > p_{\ell-1} \text{ and } i \in OPT\} \\ O_3 &= \{i | c_i \leq p_{\ell-1}\}. \end{aligned}$$

We obtain our upper bound on $\mathbb{E}[f(OPT)]$ via the following manipulation, whose first, second, and last lines follow from the fact that g is a concave function satisfying $g(0) =$

0.

$$\begin{aligned}
\mathbb{E}[f(OPT)] &= \mathbb{E}[g(|OPT|)] \leq \mathbb{E}[g(|O_1|) + g(|O_2|)] \\
&\leq g(\mathbb{E}(|O_3|)) + g(\mathbb{E}(|O_2|)) \\
&= g(ns(p_{\ell-1})) + g(\mathbb{E}(|O_2|)) \\
&\leq g(ns(p_{\ell-1})) + g\left(\frac{B}{p_{\ell-1}}\right) \\
&\leq 2g\left(\frac{B}{p_{\ell-1}}\right) \\
&\leq 2r \cdot g\left(\frac{B}{p_{\ell}}\right)
\end{aligned} \tag{3}$$

□

By Theorem 3.6, with probability at least $\frac{1}{2}$ the number of accepted offers is at least $c_5 \cdot (B/p_{\ell})$. Hence, letting S denote the set of accepted offers,

$$\mathbb{E}[f(S)] \geq \frac{1}{2}c_5 \cdot g\left(\frac{B}{p_{\ell}}\right) \geq \frac{c_5}{4r} \cdot \mathbb{E}[f(OPT)],$$

where the last inequality follows from Lemma 3.7. This completes the proof that our algorithm is constant-competitive whenever f is a nondecreasing symmetric submodular function.

4. AN $O(\log N)$ POSTED PRICE MECHANISM FOR SUBMODULAR MARKETS

In this section we present a posted price mechanism which is $O(\log n)$ -competitive for any nondecreasing submodular utility function, in the secretary model. A function $f : 2^{[n]} \rightarrow \mathbb{R}_+$ is submodular if for any subsets S, T s.t. $S \subseteq T$ we have $f(S \cup \{a\}) - f(S) \geq f(T \cup \{a\}) - f(T)$. The function is called *nondecreasing* if $S \subseteq T$ implies $f(S) \leq f(T)$.

THEOREM 4.1. *For any nondecreasing submodular procurement market there is a randomized posted price budget feasible mechanism which is universally truthful and is $O(\log n)$ -competitive.*

In proof, we present the following mechanism, and show it respects the mentioned properties. In its essence, the mechanism is quite simple: it samples approximately half of the agents, rejects them and finds the agent with the highest value in the sample. It then uses this value to guess a threshold price that can be used to decide on the allocation of the remaining agents.

A Posted Price Mechanism for Submodular Functions

With probability 1/2 do ALG1:

ALG1

1. Choose $\tau \in [0, n]$ with $Pr[\tau = i] = \binom{n}{i}/2^n$.
2. Offer $p = 0$ to the first τ agents that arrive and let $v' = \max_{\{a_i: i \leq \tau\}} f(a_i)$.
3. Choose i u.a.r from $\{0, 1, 2, \dots, \lceil \log n \rceil\}$, let $t = 2^i v'$ and $B' = B$
4. For each agent $a \in \mathcal{N} \setminus \{a_1, \dots, a_{\tau}\}$:
 - a. Offer the agent $p = \frac{B'}{t} \cdot (f(S \cup \{a\}) - f(S))$ if $B' - p \geq 0$
 - b. If a accepts, add her to S and set $B' = B' - p$.

Otherwise:

ALG2

Run Dynkin's algorithm and offer B to the winner

The mechanism runs Dynkin's algorithm with probability $1/2$ to allocate to an agent with a sufficiently high value [Dynkin 1963]. That is, with probability $1/2$ the mechanism samples the first n/e agents and then, from the remaining $(1 - 1/e)n$ agents, allocates to the first agent a' for which $f(a') \geq \operatorname{argmax}_{i \in \{1, \dots, n/e\}} f(a_i)$. In expectation over the arrival order of the agents this guarantees that $f(a') \geq (1/e) \cdot \operatorname{argmax}_{a \in \mathcal{N}} f(a)$.

Since the mechanism is a randomization over two posted price mechanisms, it is truthful and individually rational. Budget feasibility is implied from the condition in (4a) which verifies that the remaining budget B' is greater than the payment. Throughout the rest of this paper we will use $OPT(\mathcal{N}')$ to denote the optimal value over the set of agents \mathcal{N}' , \mathcal{N}_1 to denote the set of agents who are in the sample and $\mathcal{N}_2 = \mathcal{N} \setminus \mathcal{N}_1$. Let $S^* = OPT(\mathcal{N})$, $S_1^* = S^* \cap \mathcal{N}_1$ and $S_2^* = S^* \cap \mathcal{N}_2^*$.

Let $v^* = \max\{f(a), a \in \mathcal{N}\}$. Consider the case when $f(v^*) \geq f(OPT)/1024$. In this case the algorithm runs Dynkin's algorithm with probability $1/2$ and gets an approximation ratio of $2048e$. So for the rest of the section we will assume that $f(a) < f(OPT)/1024, \forall a \in \mathcal{N}$. To prove the competitive ratio of the mechanism, we will use the following lemma.

LEMMA 4.2. *When $f(a) \leq \frac{f(OPT)}{1024}, \forall a \in \mathcal{N}$ then $\min\{f(S_1^*), f(S_2^*)\} \geq \frac{OPT(\mathcal{N})}{4}$ with probability at least $p \geq \frac{9}{10}$.*

PROOF. Let $S^* = \{a_1, \dots, a_\ell\}$ be the optimal solution. Without loss of generality, assume that a_1, \dots, a_ℓ are sorted according to decreasing marginal contributions, and let w_i denote the marginal contribution of a_i .

Since the agents are assumed to arrive in a uniformly random order, and τ is chosen with a specific probability distribution we have that the sampled set \mathcal{N}_1 is a uniformly random set of \mathcal{N} . Hence each agent is in \mathcal{N}_1 with probability $1/2$ independently of other agents. Consider the random variables X_1, \dots, X_ℓ , s.t. X_i takes the value w_i with if agent i belongs to \mathcal{N}_1 and 0 otherwise. Since \mathcal{N}_1 is a uniformly random set this implies that X_i takes value w_i with probability $1/2$ and 0 with probability $1/2$. Such a trick of choosing set \mathcal{N}_1 so that X_i 's are independent was first introduced by [Kleinberg 2005]. Let $X = \sum_{i=1}^{\ell} X_i$ and $\bar{X} = f(S^*) - X$. Observe that $f(S_1^*) \geq X$, and $f(S_2^*) \geq \bar{X}$ by submodularity. To show the desired properties of X and \bar{X} we will use the Chernoff bound:

THEOREM 4.3. (Chernoff Bound) *Let X_1, \dots, X_ℓ be independent random variables where X_i take values in $[0, w_i]$ and let $\mu = \mathbb{E}[\sum_{i=1}^{\ell} X_i]$. Then, for any $\delta > 0$ we have that:*

$$Pr \left[\sum_{i=1}^{\ell} X_i \geq (1 + \delta)\mu \right] \leq \left(\frac{e^\delta}{(1+\delta)^{(1+\delta)}} \right)^{\frac{\mu}{\max_i w_i}} \quad (4)$$

$$Pr \left[\sum_{i=1}^{\ell} X_i \leq (1 - \delta)\mu \right] \leq e^{\frac{-\delta^2 \mu}{2 \cdot \max_i w_i}} \quad (5)$$

Since our assumption that $\max_i w_i \leq \frac{f(S^*)}{1024}$ and that $\mu = f(S^*)/2$ the above bound implies that:

$$Pr \left[\bar{X} \leq \frac{f(S^*)}{4} \right] = Pr \left[X \geq \frac{3f(S^*)}{4} \right] \leq 1/20 \quad (6)$$

as well as:

$$\Pr\left[\bar{X} \geq \frac{3f(S^*)}{4}\right] = \Pr\left[X \leq \frac{f(S^*)}{4}\right] \leq 1/20 \quad (7)$$

Let $W_1 = \sum_{i \in S_1^*} w_i$ and $W_2 = \sum_{i \in S_2^*} w_i$. By union bound, the above inequalities imply that with probability at least $p \geq 1 - (1/20 + 1/20) = 9/10$ we have that:

$$W_2/3 \leq W_1 \leq 3W_2 \quad (8)$$

Since $OPT = W_1 + W_2$, this implies that both W_1 and W_2 are greater than $OPT/4$ with probability p . \square

LEMMA 4.4. *When $f(a) \leq \frac{f(OPT)}{1024} \leq \frac{4f(S_2^*)}{1024}$, for any realization of t where $\frac{f(S_2^*)}{64} \leq t \leq \frac{f(S_2^*)}{2}$ ALG1 is a 128-approximation to $f(S_2^*)$.*

PROOF. Let S be the set of all agents that received the mechanism's offer. Since agents are rational and the mechanism is incentive compatible, an agent a_i which rejects the mechanism's offer p_i implies that $c_i > p_i$. Note that in order to be included in S , an agent a_i needs to respect: $c_i \leq p_i \leq B'$, where B' is the remaining budget at stage i in which a_i appears. First, consider the case where every agent $a_i \in S_2^* \setminus S$ rejects the mechanism's offer in ALG1, i.e. $c_i > p_i$. In this case we have:

$$\sum_{a \in S_2^* \setminus S} \left(f(S \cup \{a\}) - f(S) \right) = \sum_{a \in S_2^* \setminus S} \left(\frac{f(S \cup a) - f(S)}{c_i} \right) \cdot c_i \quad (9)$$

$$< \sum_{a \in S_2^* \setminus S} \left(\frac{t}{B} \right) \cdot c_i \quad (10)$$

$$< \frac{f(S_2^*)}{2B} \cdot B \quad (11)$$

$$= \frac{f(S_2^*)}{2} \quad (12)$$

where (10) is due to the fact that S_2^* is a feasible solution ($\sum_{a_i \in S_2^*} c_i \leq B$), and (11) is due to the definition of t . This inequality implies:

$$f(S_2^*) - f(S) = f((S \cup S_2^* \setminus S)) - f(S) \quad (13)$$

$$\leq \sum_{a \in S_2^* \setminus S} \left(f(S \cup a) - f(S) \right) \quad (14)$$

$$\leq \frac{f(S_2^*)}{2} \quad (15)$$

and we have that $f(S) > f(S_2^*) - f(S_2^*)/2 = f(S_2^*)/2$. Hence in this case it is a factor 2 approximation to $f(S_2^*)$.

Consider the second case where there is an agent $a_i \in S_2^* \setminus S$ for which $c_i < p_i$, this implies that at stage i , $B' < p_i$. Let a_i be the first agent in $S_2^* \setminus S$ which has this property. Now, there are two (sub)cases to consider. The first case is that $B' \leq B/2$, i.e. $\sum_{a_j \in S} p_j \geq B/2$:

$$\frac{B}{2} < \sum_{a_j \in S} p_j = \sum_{a_j \in S} \frac{B}{t} \left(f(S_j \cup \{a_j\}) - f(S_j) \right) \leq B \cdot \frac{64f(S)}{f(S_2^*)} \quad (16)$$

which implies that $f(S)$ is a 128 approximation of $f(S_2^*)$. Finally, in the case where $B' < p_i$ and $B' > B/2$, we have:

$$\frac{B}{2} < B' < p_i = \frac{B}{t} \left(f(S_j \cup \{a_i\}) - f(S_j) \right) \leq \frac{B}{t} \cdot f(a_i) \leq \frac{64B}{f(S_2^*)} \cdot f(a_i) \leq \frac{64B}{256} \quad (17)$$

Here equation 17 is a contradiction. Hence the case that $B' > B/2$ never happens. \square

There are two cases to show that the algorithm is $O(\log n)$ competitive ratio.

- (1) If $f(v^*) \geq f(OPT)/1024$ then we ALG2 is a $1024e$ approximation which is run with probability $1/2$. In this case we get a $2048e$ approximation as remarked above.
- (2) If $f(a) \leq f(OPT)/1024, \forall a \in \mathcal{N}$, we consider the approximation ratio of ALG1 which is run with probability $1/2$. With probability $1/2$ v^* is included in \mathcal{N}_1 (Event T_1) and with probability at least $9/10$ we have that $f(S_2^*) \geq f(OPT)/4$ by lemma 4.2 (Event T_2). Hence by union bound both these events (Events T_1 and T_2) happen with probability at least $4/10 = 1 - (1/10 + 1/2)$. Independently of these events t such that $f(S_2^*)/64 \leq t \leq f(S_2^*)/2$ is chosen with probability $1/O(\log(n))$ which results in an approximation ratio of 128 by lemma 4.4. Hence the final approximation ratio is $\frac{1}{2} \frac{4}{10} \frac{1}{\log(n)} \frac{1}{4} \frac{1}{128} = \frac{1}{O(\log(n))}$.

This gives us Theorem 4.1 which is the main result of this section.

5. A CONSTANT-COMPETITIVE MECHANISM IN THE BIDDING MODEL

In this section we present a bidding mechanism for nondecreasing submodular markets in the secretary model. As each agent arrives, the mechanism collects its bid and must make an irrevocable decision of whether or not the agent should be allocated and how much the agent should be rewarded. We will show the following theorem.

THEOREM 5.1. *In the bidding model, for any nondecreasing submodular utility function, there is a universally truthful budget feasible mechanism which is $O(1)$ -competitive.*

The result of the previous section showed that the gap between posted price mechanisms and bidding mechanisms in nondecreasing procurement markets in the secretary model is at most $O(\log n)$. The above theorem and theorem 3.6 which are a special cases of model considered in section 4, hint towards the possibility that the gap in theorem 4.1 can be reduced to $O(1)$.

Like the mechanism from the previous section, the mechanism here will also sample the agents and use a threshold value to decide on the allocation. In the process of estimating the threshold, the mechanism will compute an approximation of the optimal solution of the bids of the first half of the agents. A modification of the greedy algorithm that sorts agents according to their density – their marginal contribution normalized by their cost – achieves an approximation ratio of $e/(e-1)$ [Khuller et al. 1999; Sviridenko 2004] which is known to be optimal [Feige 1998]. We use $A(\mathcal{N}_1)$ to denote the value of this algorithm computed over a subset of agents $\mathcal{N}_1 \subseteq \mathcal{N}$.

An Online Bidding Mechanism for Submodular Functions

With probability $1/2$ do ALG1:

ALG1

1. Choose $\tau \in [0, n]$ with $Pr[\tau = i] = \binom{n}{i}/2^n$.
2. Let \mathcal{N}_1 be the first τ agents that arrive, $\mathcal{N}_2 = \mathcal{N} \setminus \mathcal{N}_1$, and $B' = B$.
3. For each agent $a_i \in \mathcal{N}_1$:
 - a. Add the a_i 's bid to the sample;
 - b. Reject a_i 's bid
4. Compute $t = A(\mathcal{N}_1)/8$
5. For each agent $a_i \in \mathcal{N}_2$:

If $c_i \leq p_i = \frac{B}{t} \cdot (f(S \cup \{a\}) - f(S))$ and $B' - p_i \geq 0$:
 add a_i to S , pay her p_i and set $B' = B - p_i$.

Otherwise:

ALG2

Run Dynkin's algorithm and pay the winner B .

The mechanism above is similar to the one from Section 4 though it does not guess a threshold at random, but rather leverages the access it has to agents' bids to compute a "good" threshold value t from the sample. We will show that this threshold estimates OPT well, which allows for a good approximation. First though, we verify truthfulness.

LEMMA 5.2. *The above mechanism is universally truthful.*

PROOF. Fix an arrival sequence of the agents. An agent $a \in \mathcal{N}_1$ is always rejected by the mechanism, regardless of her bid, and she therefore cannot be better off by misreporting her cost. For each agent $a \in \mathcal{N}_2$ the threshold price it is offered is independent of her bid since agents have no control over their arrival order, and cannot control their marginal contributions.

Assume for purpose of contradiction that an agent a_i benefits from declaring a bid $b_i \neq c_i$. If the agent is allocated when declaring her true cost c_i , she receives the same price as long as her bid is below the threshold. Therefore she cannot benefit from declaring a cost which is higher or lower than her true cost and is below the threshold. Bidding above the threshold excludes the agent from the allocation and her utility in this case is 0, and she is worst off since the fact that she is allocated implies that her cost is below the threshold.

In case the agent is not allocated, she cannot benefit from bidding a higher cost or any cost above the threshold. In case the agent declares a cost which is below the threshold, she can get allocated. However, since her true cost is above the threshold, her utility would be negative, and she is better off declaring her true cost. \square

Budget feasibility and individual rationality are maintained by enforcing the conditions in step (4) of the mechanism. Showing that the mechanism is indeed constant competitive would there imply the main result of this section.

LEMMA 5.3. *The mechanism is $O(1)$ -competitive.*

PROOF. Once again we have two cases to argue about. The first case consider that $f(v^*) \geq f(OPT)/1024$. Then with probability $1/2$ we are running dynkins algorithm and getting in expectation $\frac{1}{e} \max_{a \in \mathcal{N}} f(a)$. Hence this is a $2 \cdot e \cdot 1024 = 2048e$ approximation.

As in Lemma 4.2 let $S^* = \{a_1, \dots, a_\ell\}$ be the optimal solution $OPT(\mathcal{N})$ and without loss of generality assume that a_1, \dots, a_ℓ are sorted according to decreasing marginal contributions. Let w_i denote the marginal contribution of a_i , $S_1^* = S^* \cap \mathcal{N}_1$, $S_2^* = S^* \setminus S_1^*$ and $W_1 = \sum_{a_i \in S_1^*} w_i$ and $W_2 = \sum_{a_i \in S_2^*} w_i$.

Consider the second case where $\forall a \in \mathcal{N}, f(a) \leq f(OPT)/1024$. In such a case by lemma 4.2 with probability $9/10$ we have that $\min\{f(S_1^*), f(S_2^*)\} \geq \frac{OPT(\mathcal{N})}{4}$. Given this event we will show that

$$\frac{f(S_2^*)}{2} \leq t \leq \frac{f(S_2^*)}{64}; \quad (18)$$

Recall that in Lemma 4.4, we showed that using a threshold with property (18) to allocate to agents in \mathcal{N}_2 if and only if the ratio between the threshold and their marginal contribution exceeds their cost to budget ratio (as described in step (4) of the mechanism), guarantees that in expectation the set of the allocated agents is a constant factor approximation of $f(S_2^*)$. Hence, showing property (18) above would imply that with a constant probability the value of the set of agents allocated by the mechanism is a constant factor approximation of $f(S_2^*)$. From Lemma 4.2 this implies the mechanism is constant competitive.

To compute the threshold t , in step (3) we apply the greedy algorithm for submodular maximization under a budget constraint on the sample \mathcal{N}_1 . This algorithm is guaranteed to provide at least a $(1 - 1/e)$ fraction of $OPT(\mathcal{N}_1)$. Due to the decreasing marginal utilities property of the submodular function, we have that $OPT(\mathcal{N}_1) \geq f(S_1^*) \geq W_1$ and that $OPT(\mathcal{N}_2) \geq f(S_2^*) \geq W_2$. From Lemma 4.2 we know that there is a constant probability for which $OPT(\mathcal{N})/4 \leq \min\{f(S_1^*), f(S_2^*)\}$. The threshold t can be bounded from above:

$$t = \frac{A(\mathcal{N}_1)}{8} \leq \frac{OPT(\mathcal{N}_1)}{8} \leq \frac{OPT(\mathcal{N})}{8} \leq \frac{f(S_2^*)}{2} \quad (19)$$

To bound t from below we use $\gamma = e/(e - 1)$:

$$t = \frac{A(\mathcal{N}_1)}{8} \geq \frac{OPT(\mathcal{N}_1)}{8\gamma} \geq \frac{f(S_1^*)}{8\gamma} \geq \frac{f(S^*)}{32\gamma} \geq \frac{f(S_2^*)}{32\gamma} \geq \frac{f(S_2^*)}{64} \quad (20)$$

Hence by lemma 4.4 and lemma 4.2 we get that it is a $O(1)$ approximation.

□

6. DISCUSSION AND OPEN QUESTIONS

In this paper we investigated online budget-feasible mechanism design. Our main positive results are as follows. We present a constant-competitive posted price mechanism when agents are identically distributed and the buyer has a symmetric submodular utility function. For nonsymmetric submodular utilities, under the random ordering assumption we give a posted price mechanism that is $O(\log n)$ -competitive and a truthful mechanism that is $O(1)$ -competitive but uses bidding rather than posted pricing.

The main question left open by this work is whether there exists a $O(1)$ -competitive posted price mechanism in the nonsymmetric submodular case.

References

- BABAIOFF, M., BLUMROSEN, L., DUGHMI, S., AND SINGER, Y. 2011. Posting prices with unknown distributions. In *ICS*. 166–178.
- BABAIOFF, M., IMMORLICA, N., AND KLEINBERG, R. 2007. Matroids, secretary problems, and online mechanisms. In *SODA*. 434–443.

- BEI, X., CHEN, N., GRAVIN, N., AND LU, P. 2011. Budget feasible mechanism design via random sampling. *CoRR abs/1107.2994*.
- BESBES, O. AND ZEEVI, A. 2009. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research* 57, 1407–1420.
- BLUMROSEN, L. AND NISAN, N. 2005. On the computational power of iterative auctions. In *ACM Conference on Electronic Commerce*. 29–43.
- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *STOC*. 311–320.
- CHEN, N., GRAVIN, N., AND LU, P. 2011. On the approximability of budget feasible mechanisms. In *SODA*. 685–699.
- DOBZINSKI, S., PAPADIMITRIOU, C., AND SINGER, Y. 2011. Mechanisms for complement-free procurement. In *EC*.
- DYNKIN, E. B. 1963. The optimum choice of the instant for stopping a Markov process. *Soviet Math. Dokl* 4.
- FEIGE, U. 1998. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM* 45, 4, 634–652.
- KHULLER, S., MOSS, A., AND NAOR, J. S. 1999. The budgeted maximum coverage problem. *Inf. Process. Lett.* 70, 1, 39–45.
- KLEINBERG, R. D. 2005. A multiple-choice secretary algorithm with applications to online auctions. In *SODA*. 630–631.
- KLEINBERG, R. D. AND LEIGHTON, F. T. 2003. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *FOCS*. 594–605.
- SINGER, Y. 2010. Budget feasible mechanisms. In *FOCS*. 765–774.
- SVIRIDENKO, M. 2004. A note on maximizing a submodular set function subject to a knapsack constraint. *Oper. Res. Lett.* 32, 1, 41–43.