Online Prophet-Inequality Matching with Applications to Ad Allocation

Saeed Alaei, University of Maryland at College Park. saeed@cs.umd.edu
MohammadTaghi Hajiaghayi, University of Maryland at College Park. hajiagha@cs.umd.edu
Vahid Liaghat, University of Maryland at College Park. vliaghat@cs.umd.edu

We study the problem of online prophet-inequality matching in bipartite graphs. There is a static set of bidders and an online stream of items. We represent the interest of bidders in items by a weighted bipartite graph. Each bidder has a capacity, i.e., an upper bound on the number of items that can be allocated to her. The weight of a matching is the total weight of edges matched to the bidders. Upon the arrival of an item, the online algorithm should either allocate it to a bidder or discard it. The objective is to maximize the weight of the resulting matching. We consider this model in a stochastic setting where we know the distribution of the incoming items in advance. Furthermore, we allow the items to be drawn from different distributions, i.e., we may assume that the t^{th} item is drawn from distribution $\mathcal{D}_t$. In contrast to i.i.d. model, this allows us to model the change in the distribution of items throughout the time. We call this setting the *Prophet-Inequality Matching* because of the possibility of having a different distribution for each time. We generalize the classic prophet inequality by presenting an algorithm with the approximation ratio of $1 - \frac{1}{\sqrt{k+3}}$ where k is the minimum capacity. In case of $k = 2$, the algorithm gives a tight ratio of $\frac{1}{2}$ which is a different proof of the prophet inequality.

We also consider a model in which the bidders do not have a capacity, instead each bidder has a budget. The weight of a matching is the minimum of the budget of each vertex and the total weight of edges matched to it, when summed over all bidders. We show that if the bid to the budget ratio of every bidder is at most $\frac{1}{k}$ then a natural randomized online algorithm has an approximation ratio of $1 - \frac{k^k}{e^k k!} \approx 1 - \frac{1}{\sqrt{2\pi k}}$ compared to the optimal offline (in which the ratio goes to 1 as k becomes large).

We also present the applications of our model in Adword Allocation, Display Ad Allocation, and AdCell Model.

Categories and Subject Descriptors: F.2.1 [**Theory of Computation**]: Analysis of Algorithms

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Online Algorithms, Online Matching, Prophet Inequality, Ad Allocation

1. INTRODUCTION

The topic of Prophet Inequality has been studied in optimal stopping theory since the 1970s [Krengel and Sucheston 1977, 1978; Kennedy 1987] and more recently in computer science [Hajiaghayi et al. 2007]. In prophet inequality setting, given the distribution of a sequence of random variables $x_1, \cdots, x_n$ an onlooker has to choose from the succession of these values, where x_t is revealed to us at time t. The onlooker can only choose a certain number of values (called her capacity) and cannot choose a past value. The onlooker's goal is to maximize her revenue. The inequality has been interpreted as meaning that a prophet with complete foresight has only a bounded advantage over an onlooker who observes the variables one by one, and this explains the name Prophet Inequality.

The most basic Prophet Inequality discovered by Krengel, Sucheston, and Garling in 1970's concerns the case in which the values are chosen independently from known distributions (but not necessary identical) and the onlooker can only choose one value. Using a very simple example, they showed no online algorithm can be better than $\frac{1}{2}$-competitive [Krengel and Sucheston 1977]. Let

This work is supported in part by NSF CAREER Award 1053605, ONR Young Investigator Award N000141110662, and Google Faculty Research Award FA8650-11-1-7162.

EC'12, June 4–8, 2012, Valencia, Spain. Copyright 2012 ACM 978-1-4503-1415-2/12/06...$10.00.

$q = \frac{1}{\epsilon}$, and $q' = 0$. The first value, i.e., x_1 is always 1. The second value is either q with probability ϵ or q' with probability $1-\epsilon$. Observe that the expected revenue of any (randomized) online algorithm is $\max\{1, \epsilon(\frac{1}{\epsilon})\} = 1$. However the prophet, i.e., the optimum offline solution would choose q if it arrives, and he would choose the first value otherwise. Thus the optimum offline revenue is $(1-\epsilon) \times 1 + \epsilon(\frac{1}{\epsilon}) \approx 2$. We note that without considering stochastic assumptions, we cannot hope to get any constant competitive ratio since this model generalizes the secretary problem in this case. We study the concept of prophet inequality as a stochastic online bipartite matching problem where we have multiple onlookers with arbitrary capacities:

Prophet-Inequality Matching (PIM): *There is a set I of* bidders *and a set J of* item types*. Each item type is a vector of size $|I|$ in which the ith entry shows u_{ij}, the bid of bidder i for item type j. In each* time $t \in \{1, \ldots, T\}$ *an* item x_t *is revealed to the algorithm, where T is the total number of items. The type of x_t is chosen from the* distribution p_t *over all the item types. When an item arrives we have to either allocate it to a bidder or discard it right away. However, each bidder i has a capacity c_i, which means we cannot allocate more than c_i items to i. The benefit of allocating an item of type j to the bidder i is u_{ij}.*
We represent the bidders and the item types as a bipartite graph in which we put an edge of weight u_{ij} between bidder i and item type j. In our online stochastic setting, we assume that the distributions (i.e., p_t's) are independent and known and the supporting set of these distributions is finite. We also assume that the capacities have a lower bound $\mathbf{k}_{\geq 1}$ for some $\mathbf{k} \in \mathbb{N}$, i.e., $c_i \geq \mathbf{k}$ for all $i \in I$.

We note that the problem considered by Krengel et. al. [Krengel and Sucheston 1977] is a special case of this model where we have only one bidder with capacity one. We generalize the $\frac{1}{2}$-competitive ratio given in [Krengel and Sucheston 1977] by presenting an algorithm with an approximation ratio of $1 - \frac{1}{\sqrt{\mathbf{k}+3}}$ where $\mathbf{k}$ is the minimum capacity (in case of $\mathbf{k} = 1$ we get the same $\frac{1}{2}$-competitive ratio). Note that as $\mathbf{k}$ grows, we rapidly get closer to optimum. The intuition is that when $\mathbf{k}$ is large, the online algorithm has more room to make mistakes; in other words, a single non-optimal allocation only wastes $\frac{1}{\mathbf{k}}$ fraction of the capacity. Another way to put it is that the total demand for capacity, which is the sum of the individual demands, is closer to its mean when the total demand consists of a large number of small demands, as opposed to a few.

Prophet Inequality has been extensively studied in case of one bidder and mainly capacity one. However, comparatively much less is known for the case of capacity more than one. Assaf et al. [Assaf and Samuel-Cahn 2000] studied this problem when the revenue is the maximum of the assigned values. They proved there exists an online algorithm for which the expected "maximum" of the k choices is within a factor $k/(k+1)$ of the prophet's payoff. In a matching model, it is more natural to consider the expected sum of the choices. Kennedy [Kennedy 1987] considered the online expected sum compared to the expected "maximum" value. He gives a recursive formula for finding the best ratio, however it is more interesting to compare the online expected sum to the offline expected sum. Let M_k and V_k denote the expected sum of the k largest values and the expected revenue of the onlooker. For one bidder and capacity k, the goal is to design an algorithm which maximizes the ratio $\beta_k = V_k/M_k$. Hajiaghayi et al. [Hajiaghayi et al. 2007] first showed that $\beta_k \leq 1 - \mathcal{O}(\frac{\sqrt{\ln k}}{\sqrt{k}})$. See [Alaei 2011] for an improvement of this bound to $1 - \mathcal{O}(\frac{1}{\sqrt{k}})$. Alaei [Alaei 2011] gave an application of the k-choice prophet inequality to *(offline) optimal bayesian combinatorial auction design*. In this paper we generalize these results to encompass a stochastic (online) b-matching problem motivated by Internet advertising. In [Alaei 2011] an involved randomized approach (gamma-conservative magician) is used to prove the $1 - \frac{1}{\sqrt{\mathbf{k}+3}}$ approximation ratio. However, in our model we use a simple and more natural dynamic programming algorithm and we use a combination of dual fitting and the sand-barrier theorem to analyse the competitive ratio of our algorithm.

More precisely, we use a combination of a very intuitive linear programming (LP) approach and dynamic programming to achieve the desired competitive ratio. Based on the given distributions, we know how many times we expect to see an item type (which may not be an integer). The algorithm constructs an LP based on these expected values, which we call Expected LP, and uses the solution to the Expected LP as a guideline for suggesting the online items to the bidders. Then we use a relatively simple dynamic programming in combination with this LP solution to check whether we should assign the suggested item to the bidder or discard it. Usually it is not easy to combine and analyze the combination of a dynamic programming approach with an online LP-based algorithm. However, using an approach similar to "dual fitting" [Jain et al. 2003], we demonstrate an interesting analysis of this combination and prove the $1 - \frac{1}{\sqrt{\mathbf{k}+3}}$-competitive ratio for our algorithm. The approach can be of its own interest in analyzing the combination of dynamic programming and more recently LP-based algorithms for other online problems.

The closest work to the capacitated model might be of [Devanur et al. 2011]. They consider a resource allocation framework where the queries arrive i.i.d. with unknown distributions. They give a $(1-\mathcal{O}(\epsilon))$-competitive algorithm where the profit contributed by any single request to the optimal profit is at most $\mathcal{O}(\frac{\epsilon^2}{\log(n/\epsilon)})$. We note that in there resource allocation framework, each bidder has a uniform valuation over items in which the bidder is interested. They also prove the same results for the case that distribution may change by time, however the adversary is restricted to choose distributions with roughly the same expected fractional optimum value. In this paper we consider a stronger stochastic assumption, however bidders may have different bids for different items and the (known) distributions may change arbitrary over time. The competitive ratio of our algorithm solely depends on k which goes quickly to one. Indeed, if each bidder accepts at least 100 items, our online algorithm gets more than 0.9 of the revenue of the fractional offline optimum.

We also consider the prophet inequality matching with budget constraints studied in the literature [Feldman et al. 2009b; Manshadi et al. 2011; Mahdian and Yan 2011; Karande et al. 2011]. In this model we do not have capacities, instead each bidder has a maximum budget.

Budgeted Prophet-Inequality Matching (BPIM): *There is a set I of* bidders *and a set J of* item types. *The bid of bidder i for item type j is u_{ij}. At time t an item x_t arrives which is drawn from the distribution D_t over J. Upon arrival, we have to either allocate the item to a bidder or discard it. Each bidder i has a budget limit b_i. The revenue that we get from each bidder is the minimum of his budget and the sum of his bids for queries allocated to him. The objective of the problem is to design an allocation algorithm which maximizes the total revenue.*
In real world the bids are much smaller than the budgets, thus we assume for some $\mathbf{k}_{\geq 1} \in \mathbb{R}$, we have for all $i \in I$ and $j \in J$, $u_{ij} \leq \frac{b_i}{\mathbf{k}}$.

The stochastic online matching problem has been extensively studied in recent years [Feldman et al. 2009b; Manshadi et al. 2011; Mahdian and Yan 2011; Karande et al. 2011]. To the best of our knowledge, all of the existing stochastic models for this problem assume an i.i.d. distribution for the items. In our model, we allow the distribution of items to be different, i.e., for each t, we assume a possibly different distribution of the bids for the t^{th} item. The i.i.d. case can be modeled as a special case in which all items are drawn from the same distribution. The possibility of having different distributions allows us to model the change in the distribution of both items and bids throughout the day or month which makes our model more realistic (see applications in Section 1.1). We use a similar linear programming (LP) approach as in the capacitated version to design an algorithm with $1 - \frac{1}{\sqrt{2\pi\mathbf{k}}}$-competitive ratio for BPIM. Though the algorithm is simple, the analysis is relatively involved specially to show that this algorithm has a competitive ratio of $1 - \frac{1}{\sqrt{2\pi\mathbf{k}}}$.

Assuming $\mathbf{k} = 1$ (i.e., without any assumption on the ratio of bids to budgets) our algorithm is a $(1 - 1/e)$-competitive, which is non-trivial in our settings. However, the competitive ratio goes quickly to one when $\mathbf{k}$ goes to infinity. Indeed, even for $\mathbf{k} = 40$ the ratio is around 0.94. To the best of our knowledge this is the first such result for the online matching problem under the known distribution assumption with edge weights. The closest work to this result might be that of Devanor and Hayes [Devanur and Hayes 2009] who consider online weighted keyword matching under the random permutation model, which is a weaker assumption as of the known distribution model. They give a $1 - \mathcal{O}(\frac{1}{\sqrt{\mathbf{k}}})$ approximation algorithm which depends on other input variables too. Our result depends solely on $\mathbf{k}$ and holds for our model where the distributions may not be identical.

1.1. Applications to Ad Allocation and Related Work

We consider the matching version of the online prophet inequality which has direct applications in ad allocation. Selling online advertisement alongside search results is the major source of revenue for search engines like Google, Yahoo, and Bing. The problem of allocating queries to advertisers can be modeled as a generalized online matching problem [Mehta et al. 2007] in which each edge has a weight (i.e., bid of the the bidder for the query) and each vertex on the fixed part of the graph has a budget (i.e., the total budget of the bidder). The objective is to allocate queries to advertisers as they arrive online so as to maximize the weight of the matching. The weight of the matching is defined as the minimum of the budget of each bidder and the sum of the bids for queries allocated to that bidder when summed over all advertisers. While the earlier papers on this problem had mostly focused on designing algorithms with good worst-case performance in prior-free settings (e.g., [Mehta et al. 2007; Buchbinder et al. 2007]), there has been recent lines of research on designing online algorithms with good average case performance. These algorithms can be divided into two lines of research. The first line of research considers a model with unknown queries/bids but with a random order of arrival (e.g., [Goel and Mehta 2008; Devanur and Hayes 2009]). The second line of research that has attracted more attention recently considers a full stochastic setting in which the distribution of bids/queries are known in advance (e.g., [Feldman et al. 2009b; Manshadi et al. 2011; Mahdian and Yan 2011; Karande et al. 2011; Haeupler et al. 2011; Mirrokni et al. 2012]). While it was shown, that no online algorithm can approximate the offline optimal within a factor of better than $1 - \frac{1}{e}$ in the worst case instance [Mehta et al. 2007], there exist online algorithms that can approximate the offline optimal within a factor of $1 - \epsilon$ in expectation [Devanur and Hayes 2009] in the random permutation model. Recently [Mirrokni et al. 2012] gave an algorithm which is $1 - 1/e$-competitive in the adversarial model and 0.667-competitive in the stochastic model. The second line of research seems to have focused on the more restricted non-weighted model with 1-to-1 matchings, improving over the bound of $1 - \frac{1}{e}$. The best known competitive ratio in that setting is 0.703 due to [Haeupler et al. 2011]. In this paper we consider the weighted case with small bid to budget ratio in a more realistic prophet inequality setting. We show that when the bid to budget ratio goes to zero, the competitive ratio of the algorithm quickly goes to one.

In the Display Ad Allocation problem [Feldman et al. 2009a, 2010], the bidders have paid a web publisher for their ads to be shown to visitors to the website. The contract bought by bidder i specifies an integer upper bound on the number c_i of impressions that i is willing to pay for. The problem is to assign the impressions online so that while each advertiser i gets at most c_i impressions, the total weight of edges assigned is maximized. Feldman et. al. [Feldman et al. 2009a] gave a $1 - 1/e$-competitive algorithm in the adversarial model. In the stochastic setting, Display Ad Allocation can be considered a special case of PIM where all the distributions are identical.

A crucial difference between the Prophet-Inequality Matching (PIM) model and previous models considered in the literature is that the probability of the arrival of an item type may be different in different times. This allows much more flexibility in modeling the real-world online streams in a day or a month. In an extreme case the supports of distributions at different times might be completely disjoint. An example of a time sensitive environment is ad allocation in cellular networks, specially the recently proposed model *AdCell* [Alaei et al. 2011].

In AdCell model, a wireless service provider charges the advertisers for showing their ads. Each advertiser has a valuation for specific types of customers in various times and locations and has a limit on the maximum available budget. Each query is in the form of time and location and is associated with one individual customer. In order to achieve a non-intrusive delivery, only a limited number of ads can be sent to each customer. The goal is to maximize the total weight of allocated ads while respecting the budge limits of the bidders and the capacity limits of customers. Recently, new services have been introduced that offer location-based advertising over cellular network that fit in the AdCell model (e.g., ShopAlerts by AT&T). Alaei et. al. [Alaei et al. 2011] gave constant approximation algorithm for both offline and online versions of this model.

Defining different item types for a particular concept based on time allows the bidders to have different bids for a concept on different times. For example in search engines, when a user searches a keyword, the search engine may produce a different item type (based on time) and the bidders are also bidding on these item types, not on keywords. For example instead of considering an item type "restaurant", we can create different item types, "restaurant at early morning" or "restaurant at lunchtime". We adopt the distributions such that the probability of the appearance of item type "restaurant at lunchtime" is zero at time "early morning" and similarly the probability of the appearance of item type "restaurant at early morning" is zero at time "lunchtime". A fast food shop can easily make different bids for these item types since the shop has a better chance of getting a customer at lunchtime. Of course this can also be applied for weekly or monthly budgets since some bidders might be more interested in special days like weekends or Valentine day.

2. OUR RESULTS AND TECHNIQUES

In this section we formally mention our models and techniques used to design and analyze our algorithms. Let I be the set of m bidders and let J be the set of n item types. As mentioned in the introduction, we represent the bidders and the item types by a bipartite graph in which we put an edge of weight (bid) u_{ij} between bidder i and item type j. Bidder i may have budget b_i or capacity c_i. Furthermore, for a bidder i, c_i is at least $\mathbf{k}$ and in the case of budget constraint, for any item type j, $u_{ij} \leq b_i/\mathbf{k}$ for some constant $\mathbf{k}$. In each time $t \in \{1, \ldots, T\}$ an item x_t is revealed to the algorithm, where T is the total number of queries. The type of x_t is chosen from the distribution p_t over all the item types.

We can write the (offline) Budgeted Prophet-Inequality Matching (BPIM) problem and the (offline) Prophet-Inequality Matching (PIM) problem as the following integer programs (IPs) in which $\mathcal{R}_{jt}$ is the random variable indicating that the item x_t is of type j.

$$\begin{aligned}
&\text{maximize.} && \sum_i \min\Big(\sum_t \sum_j \mathbf{x}_{ijt} u_{ij}, b_i\Big) && (B)\\
&\forall j \in [n], t \in [T] && \sum_i \mathbf{x}_{ijt} \leq \mathcal{R}_{jt} && (R)\\
& && \mathbf{x}_{ijt} \in \{0,1\}
\end{aligned}$$

$$\begin{aligned}
&\text{maximize.} && \sum_i \sum_t \sum_j \mathbf{x}_{ijt} u_{ij}\\
&\forall j \in [n], t \in [T] && \sum_i \mathbf{x}_{ijt} \leq \mathcal{R}_{jt} && (R)\\
&\forall i \in [m] && \sum_t \sum_j \mathbf{x}_{ijt} \leq c_i && (C)\\
& && \mathbf{x}_{ijt} \in \{0,1\}
\end{aligned}$$

The variable $\mathbf{x}_{ijt}$ represents the event in which an item of type j is allocated to bidder i in time t. Each pair (j,t), $j \in J$ and $t \in [T]$, represents a possible item. If the item does not arrive, constraint R forces all the variables associated to that item to be zero. We note that the indicator variable $\mathcal{R}_{jt}$ is equal to one with probability $p_t(j)$. The set of constraints R ensures that each item is assigned to at most one bidder (if the item arrives), while the set of constraints C ensures that we do not assign more than c_i queries to bidder i. The answer to the first IP is $\text{OPT}_B(\mathcal{R})$, the optimum allocation for BPIM, whereas the answer to the second IP is $\text{OPT}_C(\mathcal{R})$, the optimum allocation for PIM. In the rest of the paper we may omit the indices B or C whenever the argument is valid if we choose either one as the subscript for all parameters.

The competitive ratio of an online algorithm ALG is defined as $\frac{\mathbf{E}_{\sim\mathcal{R}}[\text{ALG}(\mathcal{R})]}{\mathbf{E}_{\sim\mathcal{R}}[\text{OPT}(\mathcal{R})]}$. We analyze two intuitive algorithms to show that the competitive ratio of both algorithms go to one when $\mathbf{k}$ goes to infinity. Our main results are summarized in the following two theorems.

THEOREM 2.1. *For PIM, there is a randomized algorithm with the competitive ratio at least* $1 - \frac{1}{\sqrt{\mathbf{k}+3}}$.

THEOREM 2.2. *For BPIM, there is a randomized algorithm with the competitive ratio at least* $(1 - \frac{\mathbf{k}^{\mathbf{k}}}{e^{\mathbf{k}}\mathbf{k}!})$. *A lower bound on this ratio is* $1 - \frac{1}{\sqrt{2\pi\mathbf{k}}}$ *which goes to one as* $\mathbf{k}$ *goes to infinity.*

To be able to design the desired algorithms, we need to compare their expected revenues to the expected revenues of the optimum solutions. Hence, we need to find an upper bound for $\mathrm{E}[\text{OPT}(\mathcal{R})]$. We can get a simple, but powerful upper bound for $\mathrm{E}[\text{OPT}(\mathcal{R})]$ by solving the *Expected LP* of the aforementioned integer programs, i.e., instead of $\mathcal{R}_{jt}$ we put $\mathrm{E}[\mathcal{R}_{jt}]$, which is equal to $p_t(j)$. It can be shown that the (fractional) solution to this linear programming is indeed an upper bound on $\mathrm{E}[\mathrm{OPT}(\mathcal{R})]$ (see Lemma 3.1). Let $\mathbf{x}^*$ denote the (fractional) solution to the Expected LP. In order to get a good competitive ratio, the online algorithm uses $\mathbf{x}^*$ values to assign the queries to the bidders. Consider the situation in which in time t an item type j is revealed to us. Intuitively, $\mathbf{x}^*_{ijt}$ denotes the likelihood that the optimum solution may assign bidder i to this item. Therefore a natural choice would be to choose bidder i with probability $\frac{\mathbf{x}^*_{ijt}}{p_t(j)}$ and try to assign the item (t, j) to that bidder (or discard the item if none of the bidders were chosen). Interestingly, a very careful analysis shows that this simple and intuitive algorithm gives the competitive ratio of $1 - 1/\sqrt{2\pi\mathbf{k}}$ for BPIM and by combining it with a conditional expectation approach it gives the competitive ratio of $1 - \frac{1}{\sqrt{k+3}}$ for PIM. We note that for BPIM, we do not need to combine the LP solution with a dynamic programming since it is never good to not allocate an item.

Without any budget or capacity constraints, bidder i gets the same revenue in our algorithm as in the Expected LP, thus our only concern is to find how much a bidder lose because of the presence of these constraints. For BPIM, by simplifying the probability distributions in two steps we find the worst case of this loss ratio while in the same time, we find the example which proves the tightness of our analysis. Theorem 2.2 shows that simply by using the Expected LP, we can achieve a near one competitive ratio in many real world applications. However, with capacity constraints just using the solution to the Expected LP cannot help us. Thus we need to handle the capacities by combining the LP approach with a dynamic programming. We demonstrate an interesting analysis of this combination using a "dual fitting" approach.

Theoretically, we could design an optimal online algorithm for any online stochastic optimization problem using conditional expectation. Every time we need to make a decision, we compute the expectation of the objective value conditioned on each possible decision and then choose the one that maximizes this expectation. If the state space of the problem can be represented in a compact way, these conditional expectations can be computed efficiently using dynamic programming. However, usually this method cannot help us since the state space is exponential, as it is the case for BPIM. With capacity constraints each state should encode the remaining capacities of every bidder which leads to a huge m-dimensional state space. We break this barrier of the dynamic programming by using the Expected LP. When we use the solution to the Expected LP as a guideline we break any dependency between the bidders, hence breaking the m-dimensional table of the dynamic programming into m independent tables. This allows us to get optimal decisions when assigning an item to a bidder. However, we have to carefully analyze the loss we incur compared to the fractional solution when we decide to discard an item based on the dynamic programming table. Our analysis shows the competitive ratio of $1 - \frac{1}{\sqrt{k+3}}$ for this algorithm. We believe this approach might be of its own interest in dealing with other online LP-driven algorithms with hard constraints.

3. THE EXPECTED LP

In this section we formally define the Expected LP. We prove the solution to the Expected LP is indeed an upper bound on the expected revenue of the optimum offline solution.

DEFINITION 1 (EXPECTED LP). *Let $\mathcal{P}(b)$ be a linear programming given in Figure (i), where b is a vector of random variables. Let $\mathrm{E}[b]$ denote the expected vector corresponding to b, i.e., the ith entry of $\mathrm{E}[b]$ shows the expected value of $b(i)$. The* Expected LP *corresponding to $\mathcal{P}$, given in Figure (ii), can be obtained by replacing b by $\mathrm{E}[b]$ in $\mathcal{P}(b)$. We represent the optimum value of $\mathcal{P}(b)$ by $OPT_{\mathcal{P}(b)}$, and the optimum value of the Expected LP corresponding to $\mathcal{P}$ by $OPT^E_{\mathcal{P}}$.*

maximize.	$c^T\mathbf{x}$	maximize.	$c^T\mathbf{x}$
$s.t.$	$A\mathbf{x} \le b; \mathbf{x} \ge 0$	$s.t.$	$A\mathbf{x} \le \mathrm{E}[b]; \mathbf{x} \ge 0$
	(i)		(ii)

LEMMA 3.1. *We have $OPT^E_{\mathcal{P}} \ge \mathrm{E}\left[OPT_{\mathcal{P}(b)}\right]$.*

Proof. Let $\mathbf{x}^*(b)$ denote the optimum assignment which gives the value of $\text{OPT}_{\mathcal{P}(b)}$. Since $x^*(b)$ is a feasible solution for $\mathcal{P}(b)$, we have $A\mathbf{x}^*(b) \leq b$ and by linearity of expectation $\mathrm{E}[A\mathbf{x}^*(b)] = A\mathrm{E}[\mathbf{x}^*(b)] \leq \mathrm{E}[b]$. Thus $\mathrm{E}[\mathbf{x}^*(b)]$ is a feasible solution for the Expected LP. However we have:

$$\mathrm{E}\left[\text{OPT}_{\mathcal{P}(b)}\right] = \mathrm{E}\left[c^T\mathbf{x}^*(b)\right] = c^T\mathrm{E}[\mathbf{x}^*(b)] \leq \text{OPT}^E_{\mathcal{P}}.$$

□

We denote the Expected LPs of the relaxations of IPs for BPIM and PIM by OPT^E_B and OPT^E_C respectively. As we see in the next sections, the (fractional) solution to the Expected LP proves to be not just an upper bound but also a very powerful guideline for assigning the online queries to the bidders.

4. PROPHET-INEQUALITY MATCHING

In this section, given the set of bidders, the set of item types, and the probability distributions at each time, we give an online randomized algorithm with a competitive ratio of at least $1 - \frac{1}{\sqrt{\mathbf{k}+3}}$ for PIM (Algorithm 1), in which $\mathbf{k} = \min_i c_i$. As mentioned before, first we find the optimum solution x^* to the Expected LP which is constructed based on the given distributions (Figure 1).

$$\begin{array}{ll}
\text{maximize.} & \sum_i \sum_t \sum_j \mathbf{x}_{ijt} u_{ij} \\
\forall j \in [n], t \in [T] & \sum_i \mathbf{x}_{ijt} \leq p_j(t) \\
\forall i \in [m] & \sum_t \sum_j \mathbf{x}_{ijt} \leq c_i \\
 & \mathbf{x}_{ijt} \geq 0
\end{array}$$

Fig. 1. The Expected LP for PIM

When an item of type j arrives in time t, we randomly decide to either discard it or *reserve* it for a bidder. The decision is made by choosing a number $\alpha \in [0, 1]$ uniformly at random. The bidder i will be selected iff $\frac{\sum_{r<i} \mathbf{x}^*_{rjt}}{p_t(j)} \leq \alpha < \frac{\sum_{r\leq i} \mathbf{x}^*_{rjt}}{p_t(j)}$. The item will be discarded if $\alpha \geq \frac{\sum_{r\in[m]} \mathbf{x}^*_{rjt}}{p_t(j)}$. Since each bidder has a capacity limit, we assign the reserved item to bidder i only if this decision improves the expectation of the benefit we get from bidder i in future. Thus for each bidder we need to compute the expectation of his benefit in the future. We define $\mathcal{E}^r_{i,t}$ as the expected benefit we get from bidder i with only r remaining capacity at or after time t. As mentioned before, an interesting property of using the Expected LP as a guideline is that we can analyze the revenue of each bidder independently using a dynamic programming approach. An item (j, t) arrives with probability $p_t(j)$. In Algorithm 1 conditioned on the arrival of (j, t), we assign (j, t) to bidder i with probability $x^*_{ijt}/p_t(j)$. Thus the probability that bidder i gets the item (j, t) is x^*_{ijt}.

Now we try to find a recursive formula for $\mathcal{E}^r_{i,t}$. At time t if the algorithm does not reserve an item for bidder i, the benefit we get from bidder i would be $\mathcal{E}^r_{i,t+1}$. However, if the algorithm reserves an item of type j for bidder i we may either assign the item or discard it. The benefit we get in the former is $u_{ij} + \mathcal{E}^{r-1}_{i,t+1}$ while in the latter it is $\mathcal{E}^r_{i,t+1}$. Hence we can calculate $\mathcal{E}^r_{i,t}$ using the recursive formula below:

$$\mathcal{E}^r_{i,t} = \textstyle\sum_j \mathbf{x}^*_{ijt} \max\{u_{ij} + \mathcal{E}^{r-1}_{i,t+1}, \mathcal{E}^r_{i,t+1}\} + (1 - \sum_j \mathbf{x}^*_{ijt})\mathcal{E}^r_{i,t+1}$$

Therefore we can optimally decide if assigning the reserved item is going to improve the expectation of the benefit we get from bidder i. We note that the initial entries for bidder i are $\mathcal{E}^0_{i,t} = 0$ for any $t \in [T]$, and $\mathcal{E}^r_{i,T+1} = 0$ for any $r \in [c_i]$. Finally we charge bidder i the sum of the bids of queries assigned to him. Algorithm 1 is given in the figure.

Algorithm 1 Online Allocation for PIM

Input: The set of m bidders, the set of n item types, the distributions, $\mathbf{k}$, and the online queries.
Output: Assigning the queries to the bidders while maximizing the total benefit.
Offline Process:

Construct the Expected LP based on the given distributions (Figure 1).
Find the optimum (fractional) solution $\mathbf{x}^*$ to the Expected LP.
For each bidder i, construct a $c_i \times T$ table in which entry (r, t) contains the value of $\mathcal{E}^r_{i,t}$.

Online Scheme, Assume an item of type j is arrived at time t:

Let S_i be $\sum_{r=1}^{i} \frac{\mathbf{x}^*_{rjt}}{p_t(j)}$
Choose α uniformly in $[0, 1]$.
if $\alpha \geq S_m$ **then**
    Discard the item.
else
    Select bidder i where i satisfies $S_{i-1} \leq \alpha < S_i$.
    Let r be the remaining capacity of bidder i.
    if $u_{ij} + \mathcal{E}^{r-1}_{i,t+1} > \mathcal{E}^r_{i,t+1}$ **then**
        Assign the item to bidder i.
    else
        Discard the item.

To prove the $1 - \frac{1}{\sqrt{\mathbf{k}+3}}$-competitive ratio of the algorithm, we show that in expectation the benefit we get from each bidder i is at least $1 - \frac{1}{\sqrt{c_i+3}}\mu_i$, where μ_i is the contribution of bidder i to OPT^E_C. We note that $\mu_i = \sum_t \sum_j u_{ij}\mathbf{x}^*_{ijt}$. By linearity of expectation, we can analyze the revenue from each bidder separately. Furthermore, by definition of $\mathcal{E}^r_{i,t}$, the expected benefit we get from bidder i in algorithm 1 is $\mathcal{E}^{c_i}_{i,1}$. Therefore we need to show that $\frac{\mathcal{E}^{c_i}_{i,1}}{\mu_i}$ is at least $1 - \frac{1}{\sqrt{c_i+3}}$. In the rest of the section, we focus on an arbitrary bidder i and drop the subscript i. Furthermore, if we scale down all the bids (i.e., u_{ij}'s) by a constant factor, the ratio $\frac{\mathcal{E}^{c_i}_{i,1}}{\mu_i}$ does not change. Thus, without loss of generality, we scale all the bids such that $\mu_i = 1$. Now our goal is to find a lower bound on $\mathcal{E}^c_1$ conditioned on $\sum_t \sum_j u_j \mathbf{x}^*_{jt} = 1$. We recall that u_j is the benefit of item type j and $\mathbf{x}^*_{jt}$ is the probability that item (j, t) is going to be reserved in algorithm 1 for the bidder. We write this problem as a minimization LP and we find a lower bound for the LP using a "dual fitting" approach.

LEMMA 4.1. *Algorithm 1 gets a expected revenue of at least* $1 - \frac{1}{\sqrt{c_i+3}}\mu_i$, *where* μ_i *is the benefit from bidder* i *in* OPT^E_C.

Proof. Consider the following LP where the bids (u_j's) and the entries of the dynamic programming table are the LP variables:

$$
\begin{array}{ll}
\text{minimize} & \mathcal{E}^c_1 \\
\forall r \in [c], t \in [T] & \mathcal{E}^r_t \geq \sum_j \mathbf{x}^*_{jt} \max\{u_j + \mathcal{E}^{r-1}_{t+1}, \mathcal{E}^r_{t+1}\} + (1 - \sum_j \mathbf{x}^*_{jt})\mathcal{E}^r_{t+1} \\
\forall t \in [T] & \mathcal{E}^0_t \geq 0 \\
\forall r \in [c] & \mathcal{E}^r_{T+1} \geq 0 \\
\forall j \in [n], t \in [T] & \sum_t \sum_j u_j \mathbf{x}^*_{jt} \geq 1 \\
 & \mathcal{E}^r_t, u_j \geq 0
\end{array}
$$

The first three constraints define the entries of the dynamic programming table and the last constraint bounds the value of μ. We need to simplify this LP while ensuring that that the objective function does not increase. Thus we first narrow down the instances that would give the smallest $\mathcal{E}_1^c$. The plan of the proof is as follows. First, we show that for each t if we replace all the possible item types arriving at time t with a single item type with the probability $q_t = \sum_j \mathbf{x}^*_{jt}$ and the bid value $u_t = \sum_j u_j \frac{\mathbf{x}^*_{jt}}{q_t}$, we may only decrease $\mathcal{E}_1^c$ but μ does not change. So this replacement may only decrease $\frac{\mathcal{E}_1^c}{\mu}$, and after the replacement the size of the supporting set of the distribution at time t would be one. So without loss of generality, we only need to prove the lower bound for instances in which there is only one possible type of item at any given time. We then prove a lower bound of $1 - \frac{1}{\sqrt{c+3}}$ on the objective value of this program which implies the same lower bound for the objective value of the original program. We prove this by constructing a feasible solution for the dual of this linear program yielding a value of $1 - \frac{1}{\sqrt{c+3}}$. In what follows, we explain each step of the proof in more detail.

First assume that we replace all the possible item types arriving at time t' with a single item type with probability $q_{t'} = \sum_j \mathbf{x}^*_{jt'}$ and bid value $u_{t'} = \sum_j u_j \frac{\mathbf{x}^*_{jt'}}{q_{t'}}$. The value of μ is $\sum_t \sum_j u_j \mathbf{x}^*_{jt}$ before the replacement and it is equal to $q_{t'} u_{t'} + \sum_{t \in [T] \setminus \{t'\}} \sum_j u_j \mathbf{x}^*_{jt}$ after the replacement, thus μ is not affected. Let ${\mathcal{E}'}_t^r$ denote the expected benefit we get at or after time t with capacity r after the replacement. For all values of $t > t'$ nothing is affected so ${\mathcal{E}'}_t^r = \mathcal{E}_t^r$ for $t > t'$. Consider what happens at time t' when we make the replacement:

$$
\begin{aligned}
\mathcal{E}_{t'}^r &= \sum_j \mathbf{x}^*_{jt'} \max\{u_j + \mathcal{E}_{t'+1}^{r-1}, \mathcal{E}_{t'+1}^r\} + (1 - \sum_j \mathbf{x}^*_{jt'}) \mathcal{E}_{t'+1}^r \\
&= \sum_j \mathbf{x}^*_{jt'} \max\{u_j + \mathcal{E}_{t'+1}^{r-1}, \mathcal{E}_{t'+1}^r\} + (1 - q_{t'}) \mathcal{E}_{t'+1}^r \\
&\geq \max\{\sum_j \mathbf{x}^*_{jt'} u_j + \mathcal{E}_{t'+1}^{r-1} \sum_j \mathbf{x}^*_{jt'}, \mathcal{E}_{t'+1}^r \sum_j \mathbf{x}^*_{jt'}\} + (1 - q_{t'}) \mathcal{E}_{t'+1}^r \\
&\geq \max\{q_{t'}(u_{t'} + \mathcal{E}_{t'+1}^{r-1}), q_{t'} \mathcal{E}_{t'+1}^r\} + (1 - q_{t'}) \mathcal{E}_{t'+1}^r \\
&= {\mathcal{E}'}_{t'}^r
\end{aligned}
$$

So we proved that ${\mathcal{E}'}_{t'}^r \leq \mathcal{E}_{t'}^r$. Furthermore, we notice that for each t, $\mathcal{E}_{t-1}^r$ is an increasing function of $\mathcal{E}_t^r$ and $\mathcal{E}_t^{r-1}$ so if $\mathcal{E}_t^r$ decreases then $\mathcal{E}_{t-1}^r$ may only decrease so for all values of $t \leq t'$ we can argue that ${\mathcal{E}'}_t^r \leq \mathcal{E}_t^r$ and in particular ${\mathcal{E}'}_1^c \leq \mathcal{E}_1^c$. That means the replacement may only decrease the expected benefit of our algorithm. Therefore WLOG, it is enough to prove a lower bound for the case where for each time t there is only one item type with non-zero probability in the distribution of time t. Thus we have:

$$\mathcal{E}_t^r = q_t \max\{u_t + \mathcal{E}_{t+1}^{r-1}, \mathcal{E}_{t+1}^r\} + (1 - q_t) \mathcal{E}_{t+1}^r$$

The LP can be then simplified as follows.

$$
\begin{array}{lll}
\text{minimize:} & \mathcal{E}_1^c & \text{(LP)} \\
\forall r \in [c], \forall t \in [T]: & \mathcal{E}_t^r \geq (1 - q_t) \mathcal{E}_{t+1}^r + q_t (\mathcal{E}_{t+1}^{r-1} + u_t) & (\boldsymbol{\alpha}_t^r) \\
\forall r \in [c], \forall t \in [T]: & \mathcal{E}_t^r \geq \mathcal{E}_{t+1}^r & (\boldsymbol{\beta}_t^r) \\
 & \sum_t q_t u_t \geq 1 & (\boldsymbol{\gamma}) \\
 & u_t \geq 0, \quad \mathcal{E}_t^r \geq 0 &
\end{array}
$$

As a warm up, we first prove that the objective value of the above LP is bounded below by $\frac{1}{2}$. Later, we extend this approach to prove a lower bound of $1 - \frac{1}{\sqrt{c+3}}$ on the objective value of this LP. We can see that $\mathcal{E}_t^r$ has a decreasing marginal value in r which implies $\mathcal{E}_t^{r-1} \geq \frac{r-1}{r}\mathcal{E}_t^r$ (this can be proven by induction on t with the base case being $t = T$ and then proving for smaller t's. We prove this formally later). Combining this with the definition of $\mathcal{E}_t^r$, we get the following inequality $\mathcal{E}_t^r$:

$$\begin{aligned}
\mathcal{E}_t^r &= q_t \max(u_t + \mathcal{E}_{t+1}^{r-1}, \mathcal{E}_{t+1}^r) + (1 - q_t)\mathcal{E}_{t+1}^r \\
&= \max(q_t u_t + q_t \mathcal{E}_{t+1}^{r-1} + (1 - q_t)\mathcal{E}_{t+1}^r, \mathcal{E}_{t+1}^r) \\
&\geq \max(q_t(u_t + \frac{r-1}{r}\mathcal{E}_{t+1}^r) + (1 - q_t)\mathcal{E}_{t+1}^r, \mathcal{E}_{t+1}^r) \\
&= \max(q_t u_t + (1 - \frac{q_t}{r})\mathcal{E}_{t+1}^r, \mathcal{E}_{t+1}^r)
\end{aligned}$$

We want to minimize $\mathcal{E}_1^c$ and the argument above helps us to bound it using only the expectations with capacity c. Thus we can completely ignore the entries of table corresponding to the other capacities. Now we rewrite the previous minimization LP as the following relaxed LP with only variables $\mathcal{E}_t^c$ and u_t for $t \in [T]$. Notice that any feasible assignment for the original program is also a feasible assignment for the following program but not vice versa. So the following program is a linear relaxation of the original program and therefore its optimal value is a lower bound for the optimal value of the original program.

$$\begin{array}{lll}
\text{minimize.} & \mathcal{E}_1^c & \\
\forall t \in [T-1] & \mathcal{E}_t^c - q_t u_t - (1 - \frac{q_t}{c})E_{t+1}^c \geq 0 & (\boldsymbol{\alpha}_t) \\
 & \mathcal{E}_1^c - q_T u_T \geq 0 & (\boldsymbol{\alpha_T}) \\
\forall t \in [T-1] & \mathcal{E}_t^c - \mathcal{E}_{t+1}^c \geq 0 & (\boldsymbol{\beta}_t) \\
 & \sum_{t=1}^{T} q_t u_t \geq 1 & (\boldsymbol{\gamma}) \\
 & u_t \geq 0, \quad \mathcal{E}_t^c \geq 0 &
\end{array}$$

Next, we show that the optimal value of the above program is bounded below by $\frac{1}{2}$ which implies that the optimal value of the original program is also bounded below by $\frac{1}{2}$. To do this, we present a feasible assignment for the dual program that obtains an objective value of at least $\frac{1}{2}$. Note that the objective value of any feasible assignment for the dual program gives a lower bound on the optimal value of the primal program. The following is the dual program:

$$\begin{array}{lll}
\text{maximize.} & \boldsymbol{\gamma} & \\
\forall t \in [T] & \boldsymbol{\gamma} q_t - \boldsymbol{\alpha}_t q_t \leq 0 & (\boldsymbol{u}_t) \\
 & \boldsymbol{\alpha}_1 + \boldsymbol{\beta}_1 \leq 1 & (\boldsymbol{\mathcal{E}}_1^c) \\
\forall t \in [2 \cdots T-1] & \boldsymbol{\alpha}_t + \boldsymbol{\beta}_t - (1 - \frac{q_{t-1}}{c})\boldsymbol{\alpha}_{t-1} - \boldsymbol{\beta}_{t-1} \leq 0 & (\boldsymbol{\mathcal{E}}_t^c) \\
 & -(1 - \frac{p_{T-1}}{c})\boldsymbol{\alpha}_{T-1} - \boldsymbol{\beta}_{T-1} \leq 0 & (\boldsymbol{\mathcal{E}}_T^c) \\
 & \boldsymbol{\alpha}_t \geq 0, \quad \boldsymbol{\beta}_t \geq 0, \quad \boldsymbol{\gamma} \geq 0 &
\end{array}$$

Now, suppose we set all $\boldsymbol{\alpha}_t = \boldsymbol{\gamma}$ and $\boldsymbol{\beta}_t = \boldsymbol{\beta}_{t-1} - \frac{q_{t-1}}{c}\boldsymbol{\gamma}$ for all t except $\boldsymbol{\beta}_1 = 1 - \boldsymbol{\gamma}$. From this assignment, we get $\boldsymbol{\beta}_t = 1 - \boldsymbol{\gamma} - \boldsymbol{\gamma}\sum_{k=1}^{t-1}\frac{q_k}{c}$. Observe that we get a feasible solution as long as all $\boldsymbol{\beta}_t$'s resulting from this assignment are non-negative. Furthermore, it is easy to see that $\boldsymbol{\beta}_t > 1 - \gamma - \gamma\frac{\sum_{k=1}^{T} q_k}{c} = 1 - 2\boldsymbol{\gamma}$. Therefore, for $\boldsymbol{\gamma} = \frac{1}{2}$, all $\boldsymbol{\beta}_t$'s are non-negative and we always get a feasible solution for the dual with an objective value of $\frac{1}{2}$ which completes the main proof. Next, we present the proof of our earlier claim that $\mathcal{E}_t^{r-1} \geq \frac{r-1}{r}\mathcal{E}_t^r$.

We now prove that $\mathcal{E}_t^r \geq \frac{r}{r+1}\mathcal{E}_t^{r+1}$ by induction on t with the base case being $t = T$ which is trivially true because $\mathcal{E}_T^r = q_T u_T$ for all $r \geq 1$. Next we assume that our claim holds for $t+1$ and all values of r. We then prove it for t and all values of r as follows:

$$\begin{aligned}\mathcal{E}_t^r &= q_t \max(u_t + \mathcal{E}_{t+1}^{r-1}, \mathcal{E}_{t+1}^r) + (1-q_t)\mathcal{E}_{t+1}^r \\ &= \max(q_t(u_t + \mathcal{E}_{t+1}^{r-1}) + (1-q_t)\mathcal{E}_{t+1}^r, \mathcal{E}_{t+1}^r)\end{aligned}$$

Observe that $\max(a,b) \geq \max((1-\epsilon)a + \epsilon b, b)$ *for all* $\epsilon \in [0,1]$ *so:*

$$\begin{aligned}\mathcal{E}_t^r &\geq \max((1-\epsilon)[q_t(u_t + \mathcal{E}_{t+1}^{r-1}) + (1-q_t)\mathcal{E}_{t+1}^r] \\ &+ \epsilon\mathcal{E}_{t+1}^r, \mathcal{E}_{t+1}^r) \\ &= \max((1-\epsilon)q_t(u_t + \mathcal{E}_{t+1}^{r-1} + \frac{\epsilon}{1-\epsilon}\mathcal{E}_{t+1}^r) \\ &+ (1-q_t)\mathcal{E}_{t+1}^r, \mathcal{E}_{t+1}^r)\end{aligned}$$

Now by applying the induction hypothesis on $\mathcal{E}_{t+1}^{r-1}$ *and* $\mathcal{E}_{t+1}^r$ *and setting* $\epsilon = \frac{1}{r+1}$*:*

$$\begin{aligned}\mathcal{E}_t^r &\geq \max(\frac{r}{r+1}q_t[u_t + \frac{r-1}{r}\mathcal{E}_{t+1}^r + \frac{1}{r}\mathcal{E}_{t+1}^r] \\ &+ (1-q_t)\frac{r}{r+1}\mathcal{E}_{t+1}^{r+1}, \frac{r}{r+1}\mathcal{E}_{t+1}^{r+1}) \\ &= \frac{r}{r+1}\max(q_t[u_t + \mathcal{E}_{t+1}^r] + (1-q_t)\mathcal{E}_{t+1}^{r+1}, \mathcal{E}_{t+1}^{r+1}) \\ &= \frac{r}{r+1}\mathcal{E}_t^{r+1}\end{aligned}$$

Now, we are ready to extend the previous argument to prove that the value of (LP) is lower bounded by $1 - \frac{1}{\sqrt{c+3}}$. We do so by constructing a feasible assignment for its dual yielding an objective value of $1 - \frac{1}{\sqrt{c+3}}$. The following is the dual LP (LP).

maximize:	γ	(DLP)
$\forall t \in [T]:$	$\gamma \leq \sum_r \alpha_t^r$	(u_t)
$\forall t \in [T-1], \forall r \in [c]:$	$\alpha_{t+1}^r + \beta_{t+1}^r \leq (1-q_t)\alpha_t^r + q_t\alpha_t^{r+1} + \beta_t^r$	$(\mathcal{E}_{t+1}^r)$
	$\alpha_1^c + \beta_1^c \leq 1$	$(\mathcal{E}_1^c)$
$\forall r \neq 1:$	$\alpha_1^r + \beta_1^r \leq 0$	$(\mathcal{E}_1^r)$
	$\alpha_t^r \geq 0, \quad \beta_t^r \geq 0, \quad \gamma \geq 0$	

Before we construct a feasible assignment for the dual LP, consider the following process.

DEFINITION 2 (SAND/BARRIER PROCESS [ALAEI 2011]). *Consider a tape of infinite length with one unit of infinitely divisible sand at position* 0 *and a barrier at position* 1*. A sequence of* $q_1, \cdots, q_n$ *(with all* $q_i \in [0,1]$*) and a parameter* $\gamma \in (0,1)$ *are given as the input. The sand and the barrier are gradually moved to the right in* n *rounds. At the* i^{th} *round the following takes place. The left most* γ *fraction of the sand on the tape is selected and the* q_i *fraction of this sand is moved one position to the right. This can be defined formally as follows. Let* s_i^j *denote that amount of sand at position* j *at the beginning of round* i *and let* $y_i^j \in [0,1]$ *denote the fraction of sand selected from* *position j during round i. y_i^j are chosen such that $\sum_j s_i^j y_i^j = \gamma$ and such that for some integer θ_i, $y_i^j = 1$ for any $j < \theta_i$, and $y_i^j = 0$ for any $j > \theta_i$. So during round i, a $q_i y_i^j s_i^j$ amount of sand is moved from each position j to position $j+1$. The barrier is moved one position to the right at the end of any around when the total amount of sand at the position of the barrier is more than $1-\gamma$. Note that that the sand never crosses the barrier throughout the process. We will use λ_i to denote the position of the barrier at the beginning of round i.*

We use Def. 2 to construct a feasible assignment for (DLP) yielding an objective value of $1 - \frac{1}{\sqrt{c+3}}$. Let $\gamma = 1 - \frac{1}{\sqrt{c+3}}$ and consider the sand/barrier process corresponding to the sequence $q_1, \cdots, q_n$. Let $\alpha_i^j \leftarrow s_i^{c-j} y_i^{c-j}$ and $\beta_i^j \leftarrow s_i^{c-j}(1 - y_i^{c-j})$. Notice that γ, α_i^j and β_i^j form a feasible assignment for (DLP) yielding a value of $1 - \frac{1}{\sqrt{c+3}}$ as long as we can show that the sand does not *fall off* the LP variables. In other words, we need to show that $y_i^j = 0$ for every $j \geq c$ and every i. Equivalently, we can show that the barrier is never moved past position c.

We now prove that the barrier is never moved past position c on the tape. At the beginning of round i, let d_i denote the average distance of the sand from the origin and let d'_i denote the average distance of the sand from the barrier. Observe that $\lambda_i = d_i + d'_i$. Furthermore, notice that $d_i = \gamma q_{i-1} + d_{i-1}$, in other words, the average distance of the sand from the origin is increased exactly by γq_{i-1} during round $i-1$ (because the amount of selected sand is exactly γ and q_{i-1} fraction of the selected sand is moved one position to the right). We invoke the following theorem to bound d'_i.

THEOREM 4.2 (SAND/BARRIER [ALAEI 2011]). *Throughout the process defined in Def. 2, the average distance of the sand from the barrier is always strictly less than $\frac{1}{1-\gamma}$. In particular, at the beginning of round i this distance is strictly less than $\frac{1-\gamma^{\lambda_i}}{1-\gamma}$. The statement is true for any sequence of probabilities $q_1, \cdots, q_n$, regardless of how big $\sum_i q_i$ is.*

By applying Theorem 4.2 we get the following inequality:

$$\lambda_i = d_i + d'_i < \sum_{r=1}^{i-1} q_r \gamma + \frac{1-\gamma^{\lambda_i}}{1-\gamma} \leq k\gamma + \frac{1-\gamma^{\lambda_i}}{1-\gamma}$$

In order to show that the barrier is never moved past position c, it is enough to show that the above inequality cannot hold for $\lambda_i = k+1$. Equivalently, it is enough to prove that the following inequality always holds:

$$k + 1 \geq k\gamma + \frac{1-\gamma^{k+1}}{1-\gamma} \tag{Λ}$$

Instead of the above inequality, we can consider the stronger inequality $c + 1 \geq c\gamma + \frac{1}{1-\gamma}$ which is quadratic in γ and can be solved to get a bound of $\gamma \leq 1 - \frac{1}{1/2+\sqrt{c+1/4}}$. This bound is in fact a weaker constraint than $\gamma \leq 1 - \frac{1}{\sqrt{c+3}}$ when $c \geq 7$. It can also be verified that for $c < 7$ and $\gamma \leq 1 - \frac{1}{\sqrt{c+3}}$ the inequality (Λ) holds. So we have proved that for $\gamma \leq 1 - \frac{1}{\sqrt{c+3}}$, the barrier is never moved past position c. That completes the proof. □

Proof of Theorem 2.1: The revenue of Algorithm 1 is the sum of the benefits we get from all bidders. Now using Lemma 4.1 we have:

$$\sum_i \mathrm{E}\left[\sum_t \sum_j X_{ijt}\right] \geq \sum_i (1 - \tfrac{1}{\sqrt{c_i+3}})\mu_i = (1 - \tfrac{1}{\sqrt{\mathbf{k}+3}})\mathrm{OPT}_C^E$$

□

5. BUDGETED PROPHET-INEQUALITY MATCHING

In this section, given the set of bidders, the set of item types, and the probability distributions at each time we give an online randomized algorithm with competitive ratio of at least $1 - \frac{1}{\sqrt{2\pi\mathbf{k}}}$ (Algorithm 2). First we find the optimum (fractional) solution $\mathbf{x}^*$ to the Expected LP which is constructed based on the given distributions (Figure 2).

$$
\begin{array}{ll}
\text{maximize.} & \sum_i \sum_t \sum_j \mathbf{x}_{ijt} u_{ij} \\
\forall i \in [m] & \sum_t \sum_j \mathbf{x}_{ijt} u_{ij} \leq b_i \\
\forall j \in [n], t \in [T] & \sum_i \mathbf{x}_{ijt} \leq p_t(j); \ \mathbf{x}_{ijt} \geq 0 \\
 & \mathbf{x}_{ijt} \geq 0
\end{array}
$$

Fig. 2. The Expected LP for BPIM

Intuitively, $\mathbf{x}^*_{ijt}$ shows how many times the optimum solution has assigned the item (j,t) to bidder i in average. When an item of type j arrives in time t, we assign it to bidder i with probability $\mathbf{x}^*_{ijt}/p_t(j)$ as follows: We choose a random number $\alpha \in [0,1]$. The bidder i will be selected iff $\frac{\sum_{r<i} \mathbf{x}^*_{rjt}}{p_t(j)} \leq \alpha < \frac{\sum_{r \leq i} \mathbf{x}^*_{rjt}}{p_t(j)}$. Therefore the item will be discarded with probability $1 - \frac{\sum_{r \in [m]} \mathbf{x}^*_{rjt}}{p_t(j)}$. Finally we charge bidder i the minimum of b_i and the sum of the bids of queries assigned to i.

Algorithm 2 Online Allocation for BPIM

Input: The set of m bidders, the set of n item types, the distributions, $\mathbf{k}$, and the queries as an online stream.
Output: Assigning the queries to the bidders while maximizing the total benefit.
Offline Process:

```
Construct the Expected LP based on the given distributions (Figure 2).
Find the optimum (fractional) solution x* to the Expected LP.
```

Online Scheme, Assume an item of type j is arrived at time t:

```
Let S_i be \sum_{r=1}^{i} x*_{rjt} / p_t(j)
Choose α uniformly in [0,1].
if α ≥ S_m then
    Discard the item.
else
    Assign the item to bidder i where i satisfies S_{i-1} ≤ α < S_i.
```

An item (j,t) arrives with probability $p_t(j)$. In Algorithm 2 conditioned on the arrival of (j,t), we assign (j,t) to bidder i with probability $x^*_{ijt}/p_t(j)$. Thus the probability that bidder i gets the item (j,t) is x^*_{ijt}. Let X_{ijt} denote the random variable which shows the benefit of the bidder i from the item j in Algorithm 2, i.e.,

$$
X_{ijt} = \begin{cases} u_{ij} & \text{w.p. } \mathbf{x}^*_{ijt} \\ 0 & \text{w.p. } 1 - \mathbf{x}^*_{ijt}. \end{cases}
$$

The expected revenue of the algorithm is

$$\mathrm{E}\left[\sum_i \min\left\{\sum_t \sum_j X_{ijt}, b_i\right\}\right] = \sum_i \mathrm{E}\left[\min\left\{\sum_t \sum_j X_{ijt}, b_i\right\}\right].$$

Thus we can analyze the revenue of each bidder separately. Let μ_i denote $\mathrm{E}\left[\sum_t \sum_j X_{ijt}\right]$. By linearity of expectation $\mu_i = \sum_t \sum_j \mathbf{x}^*_{ijt} u_{ij}$ and according to the Expected LP of BPIM (Figure 2), μ_i is not greater than b_i. Furthermore, since the revenue of x^* is equal to OPT^E_B, we know that no online algorithm can do better than $\sum_i \mu_i$ in expectation. Therefore we need to find a lower bound on $\mathrm{E}\left[\min\left\{\sum_t \sum_j X_{ijt}, b_i\right\}\right]$ when μ_i is fixed.

In the rest of the section, we focus on a generic bidder i and thus we omit the i indices. For the ease of notation, let $Q \subseteq J \times [T]$ be the set of possible queries, i.e., the set of all pairs (j,t), $j \in J$ and $t \in [T]$, such that $p_t(j) > 0$. Thus instead of the variable X_{ijt}, we may simply use X_q where $q = (j,t)$. Our goal is to find a lower bound on $\mathrm{E}\left[\min\{\sum_{q \in Q} X_q, b\}\right]$ while $\mu \leq b$ is fixed. We find the worst case by simplifying the variables in two steps. Let C be the maximum value of X_q, i.e., $C = \frac{b}{\mathbf{k}}$. In the first step, we replace X_q's by independent variables with a fixed range.

LEMMA 5.1. *Let σ_q denote $\frac{\mathrm{E}[X_q]}{C}$. Consider the random variables Y_q's defined below. Replacing X_q's with Y_q's may only decrease the objective function, i.e.,* $\mathrm{E}\left[\min\{\sum_q Y_q, b\}\right] \leq \mathrm{E}\left[\min\{\sum_q X_q, b\}\right]$. *We note that the expected value of the sum of variables remains the same.*

$$Y_q = \begin{cases} C & \text{w.p. } \sigma_q \\ 0 & \text{w.p. } 1 - \sigma_q. \end{cases}$$

Proof. We replace the variables one by one and we show that in each step we may only decrease the objective function. Thus one can complete the proof using an inductive argument.
Consider a possible item q^*. We claim that replacing the variable X_{q^*} by Y_{q^*} may only decrease the objective function. It is sufficient to prove the claim when the sum of the other variables is fixed, i.e., for all $s \leq b$:

$$\mathrm{E}\left[\min\{Y_{q^*} + \sum_{q \in Q \setminus \{q^*\}} X_q, b\} | \sum_{q \in Q \setminus \{q^*\}} X_q = s\right] \leq \mathrm{E}\left[\min\{\sum_{q \in Q} X_q, b\} | \sum_{q \in Q \setminus \{q^*\}} X_q = s\right].$$

We note that the case $s \geq b$ is trivial since the expected value in this case is b regardless of the value of Y_{q^*}.

Let $r = b - s$. Since the variables are independent, we only need to show that $\mathrm{E}\left[\min\{Y_{q^*}, r\}\right] \leq \mathrm{E}\left[\min\{X_{q^*}, r\}\right]$. If $r \geq C$, the inequality clearly holds since both variables are less than C. Hence we may assume that $r < C$. Since the value of Y_{q^*} is either zero or C, the left side of the inequality is $\mathrm{E}\left[\min\{Y_{q^*}, r\}\right] = r\sigma_{q^*} = \frac{r}{C}\mathrm{E}\left[X_{q^*}\right]$. Recall that the expected values of Y_{q^*} and X_{q^*} are the same. Using Y_{q^*} we get $\frac{r}{C}$ of $\mathrm{E}\left[X_{q^*}\right]$, thus to prove the lemma we only need to show that using X_{q^*} we get at least the same ratio of $\mathrm{E}\left[X_{q^*}\right]$.

In the event $X_{q^*} \leq r$ we get the full benefit X_{q^*}, which is more than the required ratio $\frac{r}{C}$. In the event $X_{q^*} > r$ we get the benefit of r regardless of the value of X_{q^*}. Thus we have:

$$\begin{aligned} \mathrm{E}\left[\min\{X_{q^*}, r\}\right] &= \mathrm{E}\left[\min\{X_{q^*}, r\} | X_{q^*} \leq r\right] + \mathrm{E}\left[\min\{X_{q^*}, r\} | X_{q^*} > r\right] \\ &= \mathrm{E}\left[X_{q^*} | X_{q^*} \leq r\right] + \mathrm{E}\left[r | X_{q^*} > r\right] \end{aligned}$$

However, the maximum value of X_{q^*} is C and thus $r > \frac{r}{C} X_{q^*}$. By replacing r we have:

$$\begin{aligned} \mathrm{E}\left[\min\{X_{q^*}, r\}\right] &= \mathrm{E}\left[X_{q^*} | X_{q^*} \le r\right] + \mathrm{E}\left[r | X_{q^*} > r\right] \\ &\ge \mathrm{E}\left[\frac{r}{C} X_{q^*} | X_{q^*} \le r\right] + \mathrm{E}\left[\frac{r}{C} X_{q^*} | X_{q^*} > r\right] \\ &= \frac{r}{C} \mathrm{E}\left[X_{q^*}\right] \end{aligned}$$

Therefore $\mathrm{E}\left[\min\{Y_{q^*}, r\}\right] \le \mathrm{E}\left[\min\{X_{q^*}, r\}\right]$. □

Using Lemma 5.1, we can focus on the special case where variables are either zero or C. Next lemma shows that we can restrict the variables even further.

LEMMA 5.2. *For $q \in Q$, let Y_q's be independent random variables such that $Y_q = C$ w.p. σ_q; and $Y_q = 0$ w.p. $1 - \sigma_q$. Considering a fixed $\mu = \mathrm{E}\left[\sum_q Y_q\right] \le b$, the value of $\mathrm{E}\left[\min\{\sum_q Y_q, b\}\right]$ is minimized when $\sigma_q = \mu/(|Q|C)$ for all $q \in Q$.*

Proof. We prove the lemma by showing that replacing two variables of different probabilities σ_q and $\sigma_{q'}$, with two variables of the same probability $\frac{\sigma_q + \sigma_{q'}}{2}$ may only decrease the the expected revenue. We note that $\mu = C \sum_q \sigma_q$. Also recall that $b = \mathbf{k}\, C$. If all σ_q's are not equal, then we can find two queries $q', q'' \in Q$ such that $\sigma_{q'} \ne \sigma_{q''}$. We prove the lemma by showing that we may only decrease $\mathrm{E}\left[\min\{\sum_q Y_q, \mathbf{k}\, C\}\right]$ if we replace $Y_{q'}$ and $Y_{q''}$ by two copies of a similar variable $\bar{Y}$ where $\bar{Y} = C$ w.p. $\bar{\sigma} = \frac{\sigma_{q'} + \sigma_{q''}}{2}$.
Let S denote the sum of variables other than $Y_{q'}$ and $Y_{q''}$, i.e., $\sum_{q \in Q \setminus \{q', q''\}} Y_q$. In the event $S \ge \mathbf{k}C$, the values of the remaining two variables do not have any effect on the expected value. Thus

$$\begin{aligned} \mathrm{E}\left[\min\{S + Y_{q''} + Y_{q'}, \mathbf{k}\, C\} | S \ge (\mathbf{k})C\right] &= \mathrm{E}\left[\mathbf{k}\, C | S \ge (\mathbf{k})C\right] \\ &= \mathrm{E}\left[\min\{S + \bar{Y} + \bar{Y}, \mathbf{k}\, C\} | S \ge (\mathbf{k})C\right] \end{aligned}$$

In the event $S \le (\mathbf{k} - 2)C$ the replacement does not change the expectation of the sum since $\mathrm{E}\left[Y_{q'} + Y_{q''}\right] = \mathrm{E}\left[\bar{Y} + \bar{Y}\right]$. Thus

$$\begin{aligned} \mathrm{E}\left[\min\{S + Y_{q''} + Y_{q'}, \mathbf{k}\, C\} | S \le (\mathbf{k} - 2)C\right] &= \mathrm{E}\left[S + Y_{q''} + Y_{q'} | S \le (\mathbf{k} - 2)C\right] \\ &= \mathrm{E}\left[S | S \le (\mathbf{k} - 2)C\right] + C(\sigma_{q''} + \sigma_{q'}) \\ &= \mathrm{E}\left[S | S \le (\mathbf{k} - 2)C\right] + C(\bar{\sigma} + \bar{\sigma}) \\ &= \mathrm{E}\left[\min\{S + \bar{Y} + \bar{Y}, \mathbf{k}\, C\} | S \le (\mathbf{k} - 2)C\right] \end{aligned}$$

where the last two equalities hold since the variables are independent and the fact that $2\bar{\sigma} = \sigma_{q''} + \sigma_{q'}$.

In the events $S \le (\mathbf{k} - 2)C$ and $S \ge (\mathbf{k})C$, the replacement does not change the expected value. Thus we need to consider the only remaining event $S = (\mathbf{k} - 1)C$, in which we get the benefit $(\mathbf{k} - 1)C$ if both variables $Y_{q'}$ and $Y_{q''}$ are zero, and we get the benefit $\mathbf{k}C$ otherwise. Thus

$$\begin{aligned} \mathrm{E}\left[\min\{S + Y_{q''} + Y_{q'}, \mathbf{k}\, C\} | S = (\mathbf{k} - 1)C\right] &= \mathrm{E}\left[S | S = (\mathbf{k} - 1)C\right] (\mathbf{k} - \sigma_{q''}\sigma_{q'})\, C \\ &\ge \mathrm{E}\left[S | S = (\mathbf{k} - 1)C\right] (\mathbf{k} - \bar{\sigma}\bar{\sigma})\, C \\ &= \mathrm{E}\left[\min\{S + \bar{Y} + \bar{Y}, \mathbf{k}\, C\} | S = (\mathbf{k} - 1)C\right] \end{aligned}$$

Now we can simply prove the lemma by combining these three events:

$$\begin{aligned}
\mathrm{E}\left[\min\{S+Y_{q''}+Y_{q'},\mathbf{k}\,C\}\right] &= \mathrm{E}\left[\min\{S+Y_{q''}+Y_{q'},\mathbf{k}\,C\}|S\le(\mathbf{k}-2)C\right] \\
&\quad +\mathrm{E}\left[\min\{S+Y_{q''}+Y_{q'},\mathbf{k}\,C\}|S=(\mathbf{k}-1)C\right] \\
&\quad +\mathrm{E}\left[\min\{S+Y_{q''}+Y_{q'},\mathbf{k}\,C\}|S\ge(\mathbf{k})C\right] \\
&\ge \mathrm{E}\left[\min\{S+\bar{Y}+\bar{Y},\mathbf{k}\,C\}|S\le(\mathbf{k}-2)C\right] \\
&\quad +\mathrm{E}\left[\min\{S+\bar{Y}+\bar{Y},\mathbf{k}\,C\}|S=(\mathbf{k}-1)C\right] \\
&\quad +\mathrm{E}\left[\min\{S+\bar{Y}+\bar{Y},\mathbf{k}\,C\}|S\ge(\mathbf{k})C\right] \\
&= \mathrm{E}\left[\min\{S+\bar{Y}+\bar{Y},\mathbf{k}\,C\}\right]
\end{aligned}$$

Therefore $\mathrm{E}\left[\min\{\sum_q Y_q, b\}\right]$ is minimized when all the variables are the same. □

Now we claim that Algorithm 2 charges each bidder i at least $1-\frac{\mathbf{k}^{\mathbf{k}}}{e^{\mathbf{k}}\mathbf{k}!}$ of μ_i. By Lemma 5.2, we can show this by computing the lower bound when all variables are the same.

LEMMA 5.3. *For $q \in Q$, let X_q's be independent random variables in $[0, C]$. The value of $\mathrm{E}\left[\min\{\sum_q X_q, b\}\right]$ is at least $1-\frac{\mathbf{k}^{\mathbf{k}}}{e^{\mathbf{k}}\mathbf{k}!}$ of μ in expectation where $\mu = \mathrm{E}\left[\sum_q X_q\right] \le b$ and $\mathbf{k} = b/C$. This ratio is tight, i.e., for any $\epsilon > 0$ we can construct an example for some Q and $\mathbf{k}$ such that $\mathrm{E}\left[\min\{\sum_q X_q, b\}\right] \le \mu\left(1-\frac{\mathbf{k}^{\mathbf{k}}}{e^{\mathbf{k}}\mathbf{k}!}+\epsilon\right)$*

Proof. Using Lemmas 5.1 and 5.2, the value of $\mathrm{E}\left[\min\{\sum_q X_q, b\}\right]$ is minimized when

$$\forall q \in Q,\ X_q = \begin{cases} C & \text{w.p. } \sigma \\ 0 & \text{w.p. } 1-\sigma. \end{cases}$$

where $C = \frac{b}{\mathbf{k}}$ and $\sigma = \frac{\mu}{|Q|C}$.

Now we can compute the value of $\mathrm{E}\left[\min\{\sum_q X_q, b\}\right]$ in this special case.

$$\begin{aligned}
\mathrm{E}\left[\min\{\sum_q X_q,\mathbf{k}C\}\right] &= \sum_{r=1}^{\mathbf{k}} rC\binom{|Q|}{r}\sigma^r(1-\sigma)^{|Q|-r} \\
&\quad + \sum_{r=\mathbf{k}+1}^{|Q|}\mathbf{k}C\binom{|Q|}{r}\sigma^r(1-\sigma)^{|Q|-r} \ge \left(1-\frac{\mathbf{k}^{\mathbf{k}}}{e^{\mathbf{k}}\mathbf{k}!}\right)\mu
\end{aligned}$$

where $\binom{|Q|}{r}\sigma^r(1-\sigma)^{|Q|-r}$ shows the probability of the event that the values of exactly r variables are C. In what follows, we prove the last inequality which completes the proof.

First, with out loss of generality, we may assume that C is one, i.e., we scale down the bids and budgets by a factor of C. Thus now $\mathbf{k} = b$. We have:

$$\begin{aligned}
\mathrm{E}\left[\min\{\sum_q X_q,\mathbf{k}\}\right] &= \sum_{i=1}^{\mathbf{k}} i\binom{|Q|}{i}\sigma^i(1-\sigma)^{|Q|-i} \\
&\quad + \sum_{i=\mathbf{k}+1}^{|Q|}\mathbf{k}\binom{|Q|}{i}\sigma^i(1-\sigma)^{|Q|-i} \\
&= \mathbf{k} - \sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\binom{|Q|}{i}\sigma^i(1-\sigma)^{|Q|-i}
\end{aligned}$$

We need to find an upper bound for $\sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\binom{|Q|}{i}\sigma^i(1-\sigma)^{|Q|-i}$. First we replace $\binom{|Q|}{i}$ by $\frac{|Q|^i}{i!}$.

$$\begin{aligned}\sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\binom{|Q|}{i}\sigma^i(1-\sigma)^{|Q|-i} &\le \sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\frac{|Q|^i}{i!}\sigma^i(1-\sigma)^{|Q|-i}\\ &= (1-\sigma)^{|Q|-\mathbf{k}}\sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\frac{|Q|^i}{i!}\sigma^i(1-\sigma)^{\mathbf{k}-i}\end{aligned}$$

Since $\sigma=\frac{\mu}{|Q|}$, the term $(1-\sigma)^{\mathbf{k}-i}$ might be near to one in the worse case, thus in order to simplify the inequality we replace this term by one.

$$\begin{aligned}\sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\binom{|Q|}{i}\sigma^i(1-\sigma)^{|Q|-i} &\le (1-\sigma)^{|Q|-\mathbf{k}}\sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\frac{|Q|^i}{i!}\sigma^i\\ &= (1-\sigma)^{|Q|-\mathbf{k}}\sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\frac{\mu^i}{i!}\end{aligned}$$

where the last inequality is obtained by replacing σ by $\frac{\mu}{|Q|}$. Now we break the summation into two summations and simplify each one.

$$\begin{aligned}\sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\binom{|Q|}{i}\sigma^i(1-\sigma)^{|Q|-i} &\le (1-\sigma)^{|Q|-\mathbf{k}}\left[\sum_{i=0}^{\mathbf{k}}\mathbf{k}\frac{\mu^i}{i!}-\sum_{i=0}^{\mathbf{k}}i\frac{\mu^i}{i!}\right]\\ &= (1-\sigma)^{|Q|-\mathbf{k}}\left[\sum_{i=0}^{\mathbf{k}}\mathbf{k}\frac{\mu^i}{i!}-\sum_{i=0}^{\mathbf{k}-1}\frac{\mu^{i+1}}{i!}\right]\\ &= (1-\sigma)^{|Q|-\mathbf{k}}\left[\mathbf{k}\frac{\mu^{\mathbf{k}}}{\mathbf{k}!}+\sum_{i=0}^{\mathbf{k}-1}(\mathbf{k}-\mu)\frac{\mu^i}{i!}\right]\end{aligned}$$

The last equality is obtained by merging the two summations. We note that the factor $\mathbf{k}-\mu$ is independent of the summation variable. Recall that $\mathbf{k}=b$ and$b>\mu$. If $|Q|$ goes to infinity and we also continue the summation to infinity we have:

$$\begin{aligned}\sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\binom{|Q|}{i}\sigma^i(1-\sigma)^{|Q|-i} &\le \frac{1}{e^\mu}\left[\mathbf{k}\frac{\mu^{\mathbf{k}}}{\mathbf{k}!}+\sum_{i=0}^{\infty}(\mathbf{k}-\mu)\frac{\mu^i}{i!}\right]\\ &= \frac{1}{e^\mu}\left[\mathbf{k}\frac{\mu^{\mathbf{k}}}{\mathbf{k}!}+(\mathbf{k}-\mu)e^\mu\right]\end{aligned}$$

where the last equality follows from the fact that the Taylor expansion of e^x is $\sum_{n=0}^{\infty}\frac{x^i}{n!}$. Therefore an upper bound for $\sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\binom{|Q|}{i}\sigma^i(1-\sigma)^{|Q|-i}$ is $\mu\left(\frac{\mathbf{k}^{\mathbf{k}}}{e^{\mathbf{k}}\mathbf{k}!}\right)+\mathbf{k}-\mu$. The lemma follows directly from this upper bound:

$$\begin{aligned}\mathrm{E}\left[\min\{\sum_{\sigma}X_q,\mathbf{k}\}\right] &= \mathbf{k}-\sum_{i=0}^{\mathbf{k}}(\mathbf{k}-i)\binom{|Q|}{i}\sigma^i(1-\sigma)^{|Q|-i}\\ &\ge \mathbf{k}-\left[\mu\left(\frac{\mathbf{k}^{\mathbf{k}}}{e^{\mathbf{k}}\mathbf{k}!}\right)+\mathbf{k}-\mu\right]=\mu\left(1-\frac{\mathbf{k}^{\mathbf{k}}}{e^{\mathbf{k}}\mathbf{k}!}\right)\end{aligned}$$

Finally, we note that when $|Q|$ and $\mathbf{k}$ go to infinity, the above inequality is tight. Therefore for any $\epsilon > 0$ we can choose $|Q|$ and $\mathbf{k}$ large enough, such that $\mathrm{E}\left[\min\{\sum_\sigma X_q, \mathbf{k}\}\right] \le \mu\left(1 - \frac{\mathbf{k}^\mathbf{k}}{e^\mathbf{k}\mathbf{k}!} + \epsilon\right)$. □

Proof of Theorem 2.2: The revenue of Algorithm 2 is the sum of the benefits we get from all bidders. Now using Lemma 5.3 we have:

$$\sum_i \mathrm{E}\left[\min\left\{\sum_t \sum_j X_{ijt}, b_i\right\}\right] \ge \sum_i \left(1 - \frac{\mathbf{k}^\mathbf{k}}{e^\mathbf{k}\mathbf{k}!}\right)\mu_i = \left(1 - \frac{\mathbf{k}^\mathbf{k}}{e^\mathbf{k}\mathbf{k}!}\right)\mathrm{OPT}_B^E$$

□

REFERENCES

ALAEI, S. 2011. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *FOCS*.

ALAEI, S., HAJIAGHAYI, M. T., LIAGHAT, V., PEI, D., AND SAHA, B. 2011. Adcell: Ad allocation in cellular networks. In *ESA*.

ASSAF, D. AND SAMUEL-CAHN, E. 2000. Simple ratio prophet inequalities for a mortal with multiple choices. In *J. Appl. Prob.*

BUCHBINDER, N., JAIN, K., AND NAOR, J. 2007. Online primal-dual algorithms for maximizing ad-auctions revenue. In *ESA*.

DEVANUR, N. R. AND HAYES, T. P. 2009. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *EC*.

DEVANUR, N. R., JAIN, K., SIVAN, B., AND WILKENS, C. A. 2011. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *EC*.

FELDMAN, J., HENZINGER, M., KORULA, N., MIRROKNI, V. S., AND STEIN, C. 2010. Online stochastic packing applied to display ad allocation. In *ESA*.

FELDMAN, J., KORULA, N., MIRROKNI, V., MUTHUKRISHNAN, S., AND PÁL, M. 2009a. Online ad assignment with free disposal. In *WINE*.

FELDMAN, J., MEHTA, A., MIRROKNI, V. S., AND MUTHUKRISHNAN, S. 2009b. Online stochastic matching: Beating $1 - 1/e$. In *FOCS*.

GOEL, G. AND MEHTA, A. 2008. Online budgeted matching in random input models with applications to adwords. In *SODA*.

HAEUPLER, B., MIRROKNI, V. S., AND ZADIMOGHADDAM, M. 2011. Online stochastic weighted matching. In *WINE*.

HAJIAGHAYI, M. T., KLEINBERG, R. D., AND SANDHOLM, T. 2007. Automated online mechanism design and prophet inequalities. In *AAAI*.

JAIN, K., MAHDIAN, M., MARKAKIS, E., SABERI, A., AND VAZIRANI, V. V. 2003. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *J. ACM*.

KARANDE, C., MEHTA, A., AND TRIPATHI, P. 2011. Online bipartite matching with unknown distributions. In *STOC*.

KENNEDY, D. P. 1987. Prophet-type inequalities for multi-choice optimal stopping. In *Stoch. Proc. Applic.*

KRENGEL, U. AND SUCHESTON, L. 1977. Semiamarts and finite values. *Bull. Am. Math. Soc.*.

KRENGEL, U. AND SUCHESTON, L. 1978. On semiamarts, amarts, and processes with finite value. *In Kuelbs, J., ed., Probability on Banach Spaces.*.

MAHDIAN, M. AND YAN, Q. 2011. Online bipartite matching with random arrivals: An approach based on strongly factor-revealing LPs. In *STOC*.

MANSHADI, V. H., GHARAN, S. O., AND SABERI, A. 2011. Online stochastic matching: Online actions based on offline statistics. In *SODA*.

MEHTA, A., SABERI, A., VAZIRANI, U., AND VAZIRANI, V. 2007. Adwords and generalized online matching. *J. ACM*.

MIRROKNI, V. S., ZADIMOGHADDAM, M., AND GHARAN, S. O. 2012. Simultaneous approximations for adversarial and stochastic online budgeted allocation. In *SODA*.