

Online Lower Bounds via Duality

Yossi Azar *

Ilan Reuven Cohen*

Alan Roytman*

Abstract

In this paper, we exploit linear programming duality in the online setting, where input arrives on the fly, from the unique perspective of designing lower bounds (i.e., hardness results) on the competitive ratio. In particular, we provide a systematic method (as opposed to ad hoc case analysis that is typically done) for obtaining online deterministic and randomized lower bounds on the competitive ratio for a wide variety of problems. We show the usefulness of our approach by providing new, tight hardness results for three diverse online problems: the Vector Bin Packing problem, Ad-auctions (and various online matching problems), and the Capital Investment problem. Our methods are sufficiently general that they can also be used to reconstruct existing lower bounds.

Our approach is in stark contrast to previous works, which exploit linear programming duality to obtain positive results, often via the useful primal-dual scheme. We design a general recipe with the opposite aim of obtaining negative results via duality. The general idea behind our approach is to construct a parameterized family of primal linear programs based on a candidate collection of input sequences for proving the lower bound, where the objective function corresponds to optimizing the competitive ratio. Solving the parameterized family of primal linear programs optimally would yield a valid lower bound, but is a challenging task and limits the tools that can be applied, since analysis must be done precisely and exact optimality needs to be proved. To this end, we consider the corresponding parameterized family of dual linear programs and provide feasible solutions, where the objective function yields

a lower bound on the competitive ratio. This opens up additional doors for analysis, including some of the techniques we employ (e.g., continuous analysis, differential equations, etc.), as we need not be so careful about exact optimality. We are confident that our methods can be successfully applied to produce many more lower bounds for a wide array of online problems.

1 Introduction

In this work, we develop a systematic method that illustrates how to use linear programming duality to obtain online lower bounds. That is, we provide a general approach to obtain hardness results on the competitive ratio for a wide variety of problems. In contrast, previous works have mainly applied linear programming duality in the context of designing algorithms that obtain positive results on the competitive ratio. We demonstrate our approach by exhibiting new, tight online lower bounds for several different problems such as the Vector Bin Packing problem, Ad-auctions, and the Capital Investment problem. We also apply our methods to obtain improved lower bounds for various matching problems. Our approach can also be used to reconstruct existing lower bounds for a range of problems, such as online buffer management problems [14], two-sided online bipartite matching [44], various load balancing problems, and more. We are confident that the ideas we provide are capable of yielding many additional lower bounds for a rich set of online problems.

Recall that the benchmark we use to measure the performance of an online algorithm is the *competitive ratio*. In particular, we compare the objective value of an online algorithm ALG relative to the objective value of an optimal solution OPT that is omniscient and knows the entire input sequence in advance. More formally, for any input sequence σ , let $ALG(\sigma)$ denote the value of ALG on input sequence σ and $OPT(\sigma)$ denote the value of an optimal solution on the same input sequence. We say that ALG is c -competitive if $ALG(\sigma) \leq c \cdot OPT(\sigma) + a$ for any input sequence σ (where we allow some additive constant a). For randomized algorithms, the definition is the same, except that we consider $\mathbb{E}[ALG(\sigma)]$ instead of

*School of Computer Science, Tel-Aviv University. E-mails: azar@tau.ac.il, ilanrcohen@gmail.com, alan.roytman@cs.tau.ac.il. The work of all authors is partially supported by the Israel Science Foundation (grant No. 1506/16), by the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), and by the Blavatnik Fund. Alan Roytman is now at the University of Copenhagen, and is supported in part by Thorup's Advanced Grant DFF-0602-02499B from the Danish Council for Independent Research, and by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122.

$ALG(\sigma)$.

Previous works have used linear programming duality from the perspective of designing algorithms, thus obtaining positive results on the competitive ratio. In particular, the foundational primal-dual framework has been used extensively to great success in the design and analysis of algorithms for many different settings, and builds on the rich theory developed for the primal-dual framework in the offline setting. The method has been applied to design exact algorithms (e.g., computing flows in networks) along with developing approximation algorithms, with the fundamental works of [13,27]. Since then, the primal-dual framework has also played an important role in designing algorithms in the online setting, providing strong guarantees on the competitive ratio for a wide variety of problems. This includes well-known online problems such as metrical task systems [9], the k -server problem [9,10], weighted paging [8], set cover [1], and load balancing [5,38] to name a few.

In this paper, we demonstrate our approach on three online problems and show how to obtain tight lower bounds for each of them. Our systematic method can be summarized via the following recipe. Given an online problem, we first construct some parameterized collection of input sequences for the problem and then encode various constraints that any feasible algorithm (with some competitive ratio guarantee) must obey via a corresponding parameterized family of primal linear programs. We set the objective function of each primal linear program in such a way that optimizing it is equivalent to optimizing the competitive ratio. Considering the parameterized dual linear program and obtaining *any* feasible solution to it yields a valid lower bound on the competitive ratio of any algorithm.

We observe that solving the parameterized family of primal linear programs optimally would yield a valid lower bound, but is a challenging task and limits the tools that can be applied. In particular, such analysis must be done precisely and leaves no room for error, as optimal solutions must be computed. Instead, we consider producing feasible solutions to the family of dual linear programs, which still yields a lower bound on the competitive ratio. Producing feasible as opposed to optimal solutions gives us flexibility and provides some tolerance for error in analysis (while still maintaining the optimal lower bound asymptotically). This in turn enables a richer set of tools that can be applied to prove lower bounds (including the methods that we use from continuous analysis and differential equations), since we can relax the constraint of finding exactly optimal solutions.

We note that there are some previous results that use linear programming to produce online hardness results (e.g., see [43]), but mainly from the perspective of producing exact optimal solutions to primal linear programs. In addition, the notion of factor-revealing linear programs is similar in structure to our methods. Such linear programs have primarily been used to obtain positive results (i.e., algorithmic performance bounds), as in [28] for the offline setting. Factor-revealing linear programs have also been successfully used in the online setting in the context of obtaining positive results [35], along with obtaining negative results for specific algorithms [36] (as opposed to achieving negative results for the optimization problem itself).

Our techniques allow us to sidestep ad hoc case analysis (assuming we have already guessed the lower bound sequence) that is typically done to obtain online lower bounds, where the future input sequence is adaptively determined based on the algorithm's behavior in the past. In particular, we simultaneously encode and combine these cases via linear programming in such a way that producing strong online lower bounds essentially boils down to finding good solutions to linear programs. We now describe the three problems to which we apply our methods to obtain new lower bounds.

Vector Bin Packing: The first problem we consider is the classic Bin Packing problem in an online, multidimensional setting, and we refer to this problem as the Vector Bin Packing problem. In this problem, we are given vectors $v_i = (v_i(1), \dots, v_i(d)) \in [0, 1]^d$. We must partition the vectors into the minimum number of feasible sets B_1, \dots, B_m such that, for each $1 \leq j \leq m$ and each coordinate k , we have $\sum_{i \in B_j} v_i(k) \leq 1$. We refer to each set B_j as a bin, each of which has capacity 1 for each coordinate $1 \leq k \leq d$. In the online setting, d -dimensional vectors arrive in an online manner (i.e., on the fly) and must be immediately assigned to an open bin, or to a new bin, so that the capacity constraints on each bin are satisfied along every dimension. This problem has applications in cloud computing [39], where the use of huge data centers have become more prevalent, and the costs of providing power and cooling servers have skyrocketed. In fact, these costs now exceed the costs of purchasing hardware and servers [40]. Moreover, representing jobs as vectors captures the fact that resource usage is multidimensional (e.g., CPU, memory, and I/O). Assigning multidimensional jobs to machines also has applications for implementing databases for shared-nothing environments [26], and optimizing parallel queries in databases, since such tasks typically consume multiple resources [25].

In the work of [6], they focused on the setting where all vectors have small values in each coordinate relative to the size of a bin. In particular, they gave a $(1 + \epsilon)e$ -competitive algorithm for arbitrarily large d when all vectors have coordinate values smaller than $O\left(\frac{\epsilon^2}{\log d}\right)$, for any $\epsilon > 0$ (here, e denotes the base of the natural logarithm). They also defined a *splittable* model, where a vector v can be split into arbitrarily many fractions $v \cdot \alpha_1, v \cdot \alpha_2, \dots, v \cdot \alpha_k$, $\sum_i \alpha_i = 1$ (here, each $v \cdot \alpha_i$ can be placed into a different bin). In this setting, they gave an e -competitive algorithm. We show how to use our techniques to produce a matching lower bound of e in the splittable setting, which implies the same lower bound for the original problem (even when vectors are small).

Ad-auctions: The second problem we provide a tight lower bound for is the online Ad-auctions problem. In this problem, there are n bidders which are known up front. Each bidder i has a budget of $B(i)$. In addition, products arrive in an online manner (one product at a time). As each product j arrives, each buyer i provides a bid $b_{i,j}$ for the product j . The algorithm must then allocate product j to a buyer, and gains a revenue of $b_{i,j}$ (the decision of which buyer i receives item j is irrevocable). Moreover, buyers cannot be charged more than their total budget. The objective is to maximize the total revenue. In the fractional version of the problem, the algorithm may sell a fraction of each item j to multiple buyers (as long as the sum of all fractions does not exceed one), and the revenue received from each buyer is appropriately scaled according to the fraction sold. This problem was introduced in [37] and is very important for search engine companies that wish to maximize revenue. In practice, advertisers associate their advertisements with search keywords. They then place bids on the amount they pay whenever a user clicks on the advertisement. In particular, when a user submits a search query to the search engine, an ad-auction is run in which advertisements are immediately allocated to advertisers (see [42] for further motivation).

In [17], they gave a $(1 - \frac{1}{e})$ -competitive algorithm for this problem, which matches the lower bound given in [37]. The work of [17] also studied the bounded degree setting, where the number of buyers who are interested in each product is bounded by some parameter d . Under this constraint, they obtained an improved competitive ratio of $1 - (1 - \frac{1}{d})^d$. Note that this competitive ratio approaches $(1 - \frac{1}{e})$ from above for arbitrarily large d . We provide a matching lower bound of $1 - (1 - \frac{1}{d})^d$ for this prob-

lem.

Capital Investment: The third problem we consider is the online Capital Investment problem (also known as the multislope ski rental problem [3, 4, 33]), which is a generalization of the ski rental problem. In this setting, we wish to produce many units of a particular commodity at minimum cost. Over time, orders for units of the commodity arrive in an online manner, where each unit is characterized by its arrival time. In addition, we have a set of machines (which are available in advance), where each machine m_i is characterized by its capital cost c_i and its production cost p_i , and can produce arbitrarily many units of the commodity. At any time, the algorithm can choose to buy any machine for a capital cost of c_i . Moreover, the algorithm incurs a production cost of p_i if it uses machine m_i to produce one unit of the commodity. An algorithm must decide which machines to purchase and when to purchase machines, in order to minimize the total cost: the sum of capital costs plus production costs. This problem has applications in manufacturing and making investments in the future so as to minimize costs. In addition, the problem has many applications in financial settings [22] and asset allocation problems [41].

In [4], they studied the setting in which machines arrive in an online manner and can only be purchased after arrival. They gave an $O(1)$ -competitive algorithm for the problem assuming the case that lower production costs means higher capital costs. In [33], they studied the problem where all machines are available up front and gave an e -competitive algorithm for this setting. Using our techniques, we give a matching lower bound of e .

Contributions and Techniques Our main contribution is a general recipe that enables us to obtain online lower bounds via duality. Using our methods, we obtain the following lower bounds.

1. **Vector Bin Packing:** We give a tight lower bound on the competitive ratio for the Vector Bin Packing problem when vectors are small (even in the splittable, randomized setting) which approaches e for arbitrarily large d (here, d denotes the dimension of each vector). This improves upon the previous best lower bound of $\frac{4}{3}$ and matches the positive result obtained in [6].
2. **Ad-auctions:** We give a tight lower bound on the competitive ratio for the online Ad-auctions problem in the bounded degree setting, where d is an upper bound on the number of buyers who are interested in each product. We give a lower bound of $1 - (1 - \frac{1}{d})^d$ against randomized

algorithms, which improves upon the previous best result and matches the positive result given in [17].

3. **Capital Investment:** We give a tight lower bound on the competitive ratio for the online Capital Investment problem. In particular, we give a lower bound against randomized algorithms that is arbitrarily close to e , which improves over the previous best result of $\frac{e}{e-1}$ (i.e., the randomized ski rental lower bound [30]) and matches the positive result provided in [33].

Our results yield more implications.

- **Bipartite Matching:** For the online Ad-auctions problem, our lower bound also implies the same lower bound against randomized algorithms for the online matching problem [31] in the bounded degree setting, where the degree of each arriving node is bounded by d .
- **Instance-tight:** For the online Capital Investment problem, our techniques are sufficiently general that we can provide a different lower bound for each input configuration (i.e., values for machine capital costs and production costs) that is tight against each such configuration.

A few remarks regarding our results are in order, which may be of independent interest.

- **Fractional versus Randomized:** Typically, deterministic fractional lower bounds also imply lower bounds for randomized algorithms since we can view fractions as the expectation of probabilities (with the appropriate definition of a fractional algorithm). This easily holds for the online Ad-auctions problem and the online Capital Investment problem. For the Vector Bin Packing problem, the definition of an appropriate fractional algorithm is more subtle (see Appendix B), but this statement still holds.
- **Additive Term:** For the online Ad-auctions problem, an additive constant in the competitive ratio cannot reduce the competitive ratio since we can duplicate any lower bound example many times. For the online Capital Investment problem, an additive constant cannot reduce the competitive ratio since we can simply scale all costs. For the Vector Bin Packing problem, an additive constant can reduce the competitive ratio in general. However, our lower bound also holds if an additive constant is allowed. This is achieved in our lower bound since all input vectors are duplicated many times. Hence, in all of

our lower bounds, we assume that such additive constants are zero.

Recall that, for our techniques, once we construct some parameterized collection of input sequences and express constraints that any feasible algorithm must satisfy, we can run a linear program solver for some fixed value of the parameter to obtain a computational lower bound. However, we can only run the linear program solver on finitely many primal linear programs (i.e., for finitely many parameter values of our parameterized program). Hence, if we wish to obtain a lower bound for arbitrarily large parameter values (which is useful since our lower bounds improve as the parameter increases), we must provide a formal, mathematical proof. It is possible to obtain such a proof by finding an optimal solution to the parameterized primal linear program, but finding such a solution and proving it is optimal can be difficult. Note that we can always remove constraints from the primal linear program, and the optimal solution of the new linear program would still yield a valid lower bound. However, removing such constraints should be done carefully, as otherwise the negative result can become weaker.

To reconcile this issue, we consider the corresponding parameterized dual linear program. Now, obtaining *any* feasible solution to the dual linear program (which is significantly easier than finding an optimal solution to the primal linear program) yields a valid lower bound on the competitive ratio of any algorithm. As we obtain feasible dual solutions that are closer to the optimal dual solution, we improve our online lower bound.

Related Work The primal-dual framework has also been applied in the online setting in many previous works, such as online set cover, k -server, weighted paging, online mixed packing and covering, and load balancing [1, 5, 8–10, 18, 38].

There are some previous works that use linear programming as a technique to achieve online lower bounds on the competitive ratio, but mostly from the primal perspective. For instance, the classic randomized lower bound of $\frac{e}{e-1}$ for the ski rental problem can be viewed as optimally solving a parameterized primal linear program. In addition, the work of [43] also considered using the similar idea of solving a parameterized primal linear program to obtain online lower bounds for the classic Bin Packing problem. A lower bound of approximately 1.54 was given, along with weaker lower bounds when the input was constrained to have items of smaller sizes. Although these results use linear programming to prove lower bounds, the dual linear program was not considered. In particu-

lar, such approaches require solving a linear program optimally, which constrains the set of tools that can be applied in analysis.

Factor-revealing linear programs have also been used as a method to prove positive (algorithmic) bounds for problems, along with proving negative results for specific algorithms. The work of [28] studied the metric uncapacitated facility location problem in the offline setting and analyzed two greedy approximation algorithms using factor-revealing linear programs. The work of [35] studied online bipartite matching with random arrivals, and used factor-revealing linear programs to show that the online algorithm appearing in [31] is better than $(1 - \frac{1}{e})$ -competitive. In addition, the work of [36] analyzed algorithms for online matching with stochastic rewards, along with giving lower bounds for specific algorithms using factor-revealing linear programs.

Vector Bin Packing: There is a large body of work for the Vector Bin Packing problem. The offline setting for the Vector Bin Packing problem was studied in [11, 12, 19]. The online problem was studied in Azar et al. [7]. They gave a lower bound of $\Omega(d^{\frac{1}{B}-\epsilon})$ on the competitive ratio of any algorithm (for any $\epsilon > 0$), and designed an online algorithm that is $O(d^{\frac{1}{B-1}}(\log d)^{\frac{B}{B-1}})$ -competitive for any integer $B \geq 2$ (here, B denotes the ratio between the largest coordinate and each bin’s capacity). In [6], they studied the same problem where vectors are small relative to each bin’s capacity. For arbitrarily large d , they gave a $(1 + \epsilon)e$ -competitive algorithm whenever each vector’s coordinates are at most $O(\frac{\epsilon^2}{\log d})$, for any $\epsilon > 0$. In the splittable model, they gave an e -competitive algorithm for arbitrarily large d . For the single dimensional setting, where $d = 1$ (i.e., Bin Packing), there is a vast body of literature. For surveys, see [20, 24].

Ad-auctions: The online Ad-auctions problem has also received a great deal of attention in the community. It was first introduced by [37], where they gave a deterministic $(1 - \frac{1}{e})$ -competitive algorithm for the problem, under the assumption that each bidder’s budget is large relative to their bids. They built on the works in online bipartite matching [31] and online b -matching [29] (which is a special case of the online Ad-auctions problem, where all bidders have the same budget b and all bids are 0 or 1). The work of [17] gave a generalized $(1 - \frac{1}{e})$ -competitive algorithm for the problem within the primal-dual framework, which matches the bounds given in [37]. In addition, they gave a deterministic $(1 - \frac{1}{e})$ -competitive fractional algorithm for the on-

line matching problem in bipartite graphs. They also considered the bounded degree setting, in which they gave a $(1 - (1 - \frac{1}{d})^d)$ -competitive algorithm (where d denotes an upper bound on the total number of buyers who are interested in each product). There are many other works, in both the offline and online settings, that have considered maximizing the revenue of a seller in various models [2, 15, 16, 34].

Capital Investment: The online Capital Investment problem (sometimes referred to as the multislope ski rental problem [33]) and its variants have also been studied extensively in the literature. In [21], they gave a 4-competitive deterministic algorithm and a 3.618 lower bound, along with a 2.88-competitive randomized algorithm. In [45], they studied a similar model where machines have some duration and can expire. They gave a 4-competitive algorithm for this problem, along with a deterministic matching lower bound. In [33], they studied the problem where machines are not necessarily bought from scratch. They provided an e -competitive algorithm for this setting. Under an additive assumption regarding machine capital costs, they gave an improved algorithm with a competitive ratio of $\frac{e}{e-1}$. In [32], they studied the case when there are only two machines and gave matching upper and lower bounds on the competitive ratio for deterministic and randomized algorithms. In [4], they gave an $O(1)$ -competitive algorithm for the problem where machines arrive online, assuming the case that lower production costs implies higher capital costs. When both capital and production costs drop, logarithmic bounds are necessary and sufficient. The online mortgage problem has also been studied [23], which is similar to the online Capital Investment problem except that future demand is known and capital costs are fixed.

2 Multidimensional Vector Bin Packing

In this section, we study the online d -dimensional Vector Bin Packing problem in the splittable setting. In this problem, we are given vectors $\{v_1, \dots, v_n\}$ that arrive in an online manner, where $v_i = (v_i(1), \dots, v_i(d)) \in [0, 1]^d$ for all $i \in [n]$ (recall that $[n] = \{1, \dots, n\}$). We must assign incoming vectors into bins B_j such that, for each bin B_j and each coordinate k , we have $\sum_{i \in B_j} v_i(k) \leq 1$. The goal is to minimize the number of bins opened to feasibly pack all vectors. In the splittable model, a vector v can be split into arbitrarily many fractions, $v \cdot \alpha_1, v \cdot \alpha_2, \dots, v \cdot \alpha_k$, $\sum_i \alpha_i = 1$ (here, each $v \cdot \alpha_i$ can be placed into a different bin). In the splittable model, OPT is easy to compute and is given by

$OPT = \max_k [\sum_i v_i(k)]$ (see [6]).

We give a lower bound that is arbitrarily close to ϵ for the Vector Bin Packing problem (the lower bound approaches ϵ for arbitrarily large d). The lower bound we present in this section is against deterministic fractional algorithms, which corresponds to the splittable model. This implies the same lower bound for the original problem (i.e., the non-splittable setting). In particular, the value of the optimal solution in the non-splittable model can be made arbitrarily close to the optimal solution in the splittable setting if vectors can be made arbitrarily small. This can be done by scaling large vectors appropriately many times so that the maximum coordinate is at most ϵ (for any $\epsilon > 0$). We show how to adapt our lower bound techniques to obtain a lower bound against randomized algorithms in Appendix B. Before applying our technique, we prove the following useful claim.

CLAIM 1. *For any function $f(x)$ such that $f'(x)$ is a monotone non-increasing function, the following holds for all integers $i \geq j$:*

$$\sum_{r=j}^{i-1} f'(r+1) \leq f(i) - f(j) \leq \sum_{r=j}^{i-1} f'(r).$$

Proof. We observe that $f(i) - f(j) = \int_j^i f'(x) dx = \sum_{r=j}^{i-1} \int_r^{r+1} f'(x) dx$ and $f'(r+1) \leq \int_r^{r+1} f'(x) dx \leq f'(r)$, since $f'(x)$ is monotone non-increasing.

Our proof of the lower bound holds against any algorithm in the splittable setting that can open fractions of bins (i.e., where the capacity constraints are appropriately reduced).

THEOREM 2.1. *There is no deterministic fractional algorithm with a competitive ratio strictly better than ϵ for the d -dimensional Vector Bin Packing problem, where d is arbitrarily large.*

Proof.

Sequence definition: We give a sequence $\sigma(d)$ for which the competitive ratio approaches ϵ for any algorithm as d increases. For a fixed d , the sequence consists of d phases. In each phase i , for a large value A , vectors of type v_i arrive as follows: A vectors $v_1 = (1, 0, 0, \dots, 0)$, $2A$ vectors $v_2 = (\frac{1}{2}, 1, 0, \dots, 0)$, $3A$ vectors $v_3 = (\frac{1}{3}, \frac{1}{3}, 1, \dots, 0)$, \dots , dA vectors $v_d = (\frac{1}{d}, \frac{1}{d}, \frac{1}{d}, \dots, 1)$.

Clearly, OPT is $j \cdot A$ after the j^{th} phase. Therefore, any c -competitive algorithm must open at most a total amount of $c \cdot j \cdot A$ fractional bins (i.e., total sum of fractions) by the end of the j^{th} phase. For any

deterministic algorithm, we can assume without loss of generality that an online algorithm opens exactly a $c \cdot j \cdot A$ fractional amount of bins. Therefore, in each phase j , the algorithm opens a $c \cdot A$ fractional amount of additional bins. Note that we only analyze the algorithm's behavior at the end of each phase.

Primal variables: Let $x_{i,j}$ be the total fraction of vectors of type v_i that the online algorithm assigns to bins that are opened in phase j ($i \geq j$). Let c be a variable representing the competitive ratio guarantee. When writing our linear programs, we refer to each constraint by its associated dual variable.

The primal linear program:

$$\begin{aligned} \min \quad & c \\ \text{s.t.} \quad & \sum_{r=j}^d v_r(k) x_{r,j} \leq c \quad \text{cons. } z_{k,j} \\ & \sum_{r=1}^i x_{i,r} = 1 \quad \text{cons. } y_i \\ & c, x_{i,j} \geq 0, \end{aligned} \quad ,$$

where the constraint $z_{k,j}$ is $\forall k, j \in [d]$, y_i is $\forall i \in [d]$, and the nonnegativity constraints are $\forall i \geq j : i, j \in [d]$ (there is no variable $x_{i,j}$ for $i < j$). The constraint $z_{k,j}$ corresponds to the volume constraint of coordinate k in bins opened during phase j (note that the factor of A is canceled). In particular, for any dimension k and for any bins opened in phase j , the total volume that can be placed on these bins is at most $c \cdot A$ since this is how much space is available. The constraint corresponding to y_i says that all vectors of type v_i must be fully assigned to bins opened in phases 1 through i .

We omit constraints $z_{k,j}$ for $k < j$ (recall that we can always remove constraints from the primal linear program). In particular, such constraints are never tight. In addition, in constraint $z_{k,j}$ we omit the term $v_r(k) x_{r,j}$ for $r < k$ since $v_r(k) = 0$. We get $\forall k \geq j$ (for $k, j \in [d]$):

$$k \cdot x_{k,j} + \sum_{r=k+1}^d x_{r,j} \leq c \quad \text{cons. } z_{k,j}.$$

The dual linear program:

$$\begin{aligned}
& \max \sum_{r=1}^d y_r \\
& \text{s.t.: } \sum_{k=1}^d \sum_{j=1}^d z_{k,j} \leq 1 \quad \text{cons. } c \\
& y_i \leq i \cdot z_{i,j} + \sum_{r=j}^{i-1} z_{r,j} \quad \text{cons. } x_{i,j} \\
& z_{k,j} \geq 0,
\end{aligned}$$

where the constraint $x_{i,j}$ is $\forall i \geq j : i, j \in [d]$, and $z_{k,j}$ is defined $\forall k \geq j : k, j \in [d]$.

We omit constraint c by normalizing all variables by the term $\sum_k \sum_j z_{k,j}$, which ensures that the constraint corresponding to c is feasible, but modifies the goal function to $\frac{\sum_r y_r}{\sum_k \sum_j z_{k,j}}$.

Intuition for the dual variables assignment:

A natural assignment for y_i is $y_i = 1$. This assignment already yields an improved lower bound of 2 (see Appendix A for more details). We use a more sophisticated assignment for our final solution to obtain our tight lower bound given by $y_i = \frac{1}{i}$. To find an assignment for variables $z_{k,j}$, we look at the constraints corresponding to $x_{i,j}$ as if they were ‘continuous’ (i.e., for large $i \geq j$). In doing so, we consider a differentiable function $f_j(x)$, where $f'_j(k)$ approximately represents the variable $z_{k,j}$. With this interpretation, the constraint corresponding to $x_{i,j}$ can be viewed as:

$$\begin{aligned}
i \cdot f'_j(i) + \int_j^i f'_j(x) dx &\geq \frac{1}{i} \iff \\
x \cdot f'_j(x) + f_j(x) - f_j(j) &\geq \frac{1}{x}.
\end{aligned}$$

By solving this differential equation (assuming equality) with the boundary condition $f_j(j) = 0$ (since the variables $z_{k,j}$ only exist for $k \geq j$), we get

$$f_j(x) = \frac{\ln(x/j)}{x}, \quad f'_j(x) = \frac{1 - \ln(x/j)}{x^2}.$$

Formal feasible dual variables assignment:

$$y_i = \frac{1}{i}, \quad z_{k,j} = \begin{cases} \frac{1 - \ln(k/j)}{k^2} & \text{if } j \leq k \leq \lfloor e \cdot j \rfloor, \\ 0 & \text{otherwise.} \end{cases}$$

With this assignment, we need to verify that constraint $x_{i,j}$ is feasible. For $i \leq \lfloor e \cdot j \rfloor$, we get

$$\begin{aligned}
i \cdot \frac{1 - \ln(i/j)}{i^2} + \sum_{r=j}^{i-1} \frac{1 - \ln(r/j)}{r^2} &\geq \frac{1}{i} \iff \\
\sum_{r=j}^{i-1} \frac{1 - \ln(r/j)}{r^2} &\geq \frac{\ln(i/j)}{i},
\end{aligned}$$

which holds due to Claim 1 (we apply the claim with $f(x) = \frac{\ln(x/j)}{x}$). For $i > \lfloor e \cdot j \rfloor$, we get

$$\sum_{r=j}^{\lfloor e \cdot j \rfloor} \frac{1 - \ln(r/j)}{r^2} \geq \frac{1}{e \cdot j} \geq \frac{1}{i},$$

where the first inequality is by Claim 1.

Evaluating the goal function: Recall that $H(d) = 1 + \frac{1}{2} + \dots + \frac{1}{d}$ denotes the d^{th} harmonic number. Applying Claim 1, we have $\sum_{r=j}^{\lfloor e \cdot j \rfloor} z_{r,j} \leq \frac{1}{e \cdot j} + \frac{1}{j^2}$. This yields

$$\frac{\sum_{i=1}^d y_i}{\sum_{k=1}^d \sum_{j=1}^d z_{k,j}} \geq \frac{H(d)}{\frac{H(d)}{e} + \sum_j \frac{1}{j^2}} \rightarrow e,$$

since $\sum_{j=1}^d \frac{1}{j^2}$ is bounded by a constant, and $H(d) \rightarrow \infty$ as $d \rightarrow \infty$.

3 Online Ad-auctions

In the d -bounded online Ad-auctions problem, there are n bidders which are known up front, each with a budget of $B(i)$. Products arrive online, and for each product j , at most d bidders are interested in buying the product. Each such interested buyer i bids $b_{i,j}$ for product j . The mechanism then allocates product j to a buyer, and gains a revenue of $b_{i,j}$ (a buyer cannot be charged more than their budget). The objective is to maximize the total revenue. In the fractional version of the problem, the algorithm may sell fractions of each item j to multiple buyers.

We give a randomized lower bound of $1 - (1 - \frac{1}{d})^d$ for the d -bounded online Ad-auctions problem (which approaches $1 - \frac{1}{e}$). The lower bound we present in this section can also be applied to achieve the same lower bound against randomized algorithms for the online matching problem [31] where arriving nodes have a degree of at most d .

THEOREM 3.1. *There is no randomized algorithm with a competitive ratio strictly better than $1 - (1 - \frac{1}{d})^d$ for the d -bounded online Ad-auctions problem.*

Proof.

Sequence definition: Let $n = d^{d-1}$ be the number of initial bidders. Each bidder i has a budget of $B(i) = 1$. Moreover, for every bidder i and product j that i is interested in, we have $b_{i,j} = 1$. Our sequence is composed of $d - 1$ phases, in addition to a final phase. In each phase $k \in \{1, \dots, d - 1\}$, the

adversary only sells to some number of bidders R_k . In particular, the adversary introduces R_k/d products to R_k bidders by grouping the R_k bidders into R_k/d groups (each of size d) and introduces a product for each group. A bidder is in the set $[R_k]$ (recall that $[n] = \{1, \dots, n\}$) if they were in the set $[R_{k-1}]$ and did not have the highest leftover budget in their group from phase $k-1$ (in other words, from phase $k-1$ to k , the adversary drops a buyer from each group). We assume that the bidders are reindexed at the beginning of each phase, so that each bidder's index is in the set $[R_k]$. Note that $R_k = n \cdot \left(\frac{d-1}{d}\right)^{k-1}$ (since $R_1 = n$). In the last phase, the adversary introduces a single product to each of the remaining R_d players. Clearly, OPT is n since it can sell a product to each omitted bidder along with the rest of the bidders in the final phase.

Primal variables: Let $x_{k,i}$ be the amount sold to the i^{th} player in the k^{th} phase ($i \in R_k$). Let $t_{k,i}$ be the total amount sold to the i^{th} player up to the k^{th} phase ($i \in R_{k+1}$). Note that the i^{th} player in the k^{th} iteration might not be the i^{th} player in a different iteration k' (due to reindexing). We associate two variables with each player i in phase k : $x_{k,i}$ and $t_{k-1,i}$.

Let $\hat{i} = \lceil \frac{i}{d} \rceil$, and $G_a = \{(a-1) \cdot d + 1, \dots, a \cdot d - 1\}$ for all $a \in [R_k/d]$ (for a fixed phase k). Note that the set $G_a \cup \{a \cdot d\}$ represents a group of d players (namely, the players in group a are those who are interested in a particular product). We exclude player $a \cdot d$ from the set G_a for notational convenience. Recall that the notation $a \mid b$ means that a divides b . When writing our linear programs, we refer to each constraint by its associated dual variable.

The primal linear program:

$$\begin{aligned} \max \quad & \sum_{k=1}^{d-1} \sum_{i=1}^{R_k} x_{k,i} + \sum_{i=1}^{R_d} (1 - t_{d-1,i}) \\ \text{s.t.} \quad & \sum_{i \in G_a \cup \{a \cdot d\}} x_{k,i} \leq 1 && \text{cons. } y_{k,a} \\ & x_{k,d \cdot \hat{i}} + t_{k-1,d \cdot \hat{i}} \leq x_{k,i} + t_{k-1,i} && \text{cons. } w_{k,i} \\ & \sum_{i \in R_k, d \nmid i} (x_{k,i} + t_{k-1,i}) = \sum_{i \in R_{k+1}} t_{k,i} && \text{cons. } z_k \\ & x_{k,i}, t_{k,i} \geq 0, \end{aligned}$$

where the constraint $y_{k,a}$ is $\forall k \in [d-1], a \in [R_k/d]$, $w_{k,i}$ is $\forall k \in [d-1], i \in [R_k], d \nmid i$, z_k is $\forall k \in [d-1], i \in [R_k], i' \in [R_{k+1}]$. The constraint $y_{k,a}$ captures the fact that we can sell at most one product per group (one for each group a in phase k). The constraint $w_{k,i}$ identifies the bidder with the highest leftover

budget in their group during phase k . Lastly, the constraint z_k corresponds to reordering the bidders. In fact, we even allow bidders' leftover budgets to be redistributed. The goal function is the total amount sold from phases 1 to $d-1$ plus the remaining leftover budgets from bidders in R_d . Note that we define $t_{0,i} = 0$ for all i .

The dual linear program:

$$\begin{aligned} \min \quad & \sum_{k=1}^{d-1} \sum_{a=1}^{R_k/d} y_{k,a} + R_d \\ \text{s.t.} \quad & y_{k,\hat{i}} - w_{k,i} + z_k \geq 1 && \text{cons. } x_{k,i} \\ & y_{k,\hat{i}} + \sum_{r \in G_{\hat{i}}} w_{k,r} \geq 1 && \text{cons. } x_{k,i} \\ & -w_{k+1,i} - z_k + z_{k+1} \geq 0 && \text{cons. } t_{k,i} \\ & \sum_{r \in G_{\hat{i}}} w_{k+1,r} - z_k \geq 0 && \text{cons. } t_{k,i} \\ & -z_{d-1} \geq -1 && \text{cons. } t_{d-1,i} \\ & y_{k,\hat{i}}, w_{k,i} \geq 0, \end{aligned}$$

where the first set of constraints for $x_{k,i}$ ($\forall k \in [d-1], i \in [R_k], d \nmid i$) corresponds to bidders who continue on to the next phase. The second set of constraints for $x_{k,i}$ ($\forall k \in [d-1], i \in [R_k], d \mid i$) corresponds to bidders who are dropped and do not continue on. The first set of constraints for $t_{k,i}$ ($\forall k \in [d-2], i \in [R_{k+1}], d \nmid i$) and the second set of constraints for $t_{k,i}$ ($\forall k \in [d-2], i \in [R_{k+1}], d \mid i$) correspond in a similar manner as the $x_{k,i}$ constraints. The nonnegativity constraints are $\forall k \in [d-1], i \in [R_k], d \nmid i$. Note that the constraint corresponding to $t_{d-1,i}$ ($\forall i \in [R_d]$) is independent of i and appears $d-1$ times in the dual linear program.

Intuition for the dual variables assignment:

By symmetry, we assume that $w_{k,i} = w_{k,i'}$ for all i, i' and denote this common value by w_k . Similarly, we assume that $y_{k,a} = y_k$ for all a . In addition, we assume that all constraints are tight $\forall k \in [d-1]$, which yields:

$$\begin{aligned} (3.1) \quad & y_k - w_k + z_k = 1, \\ (3.2) \quad & y_k + (d-1)w_k = 1, \\ (3.3) \quad & -w_{k+1} - z_k + z_{k+1} = 0, \\ (3.4) \quad & (d-1)w_{k+1} - z_k = 0, \\ (3.5) \quad & z_{d-1} = 1. \end{aligned}$$

Formal feasible dual variables assignment:

From Equations (3.3) and (3.4), we derive $z_{k+1} = \left(\frac{d}{d-1}\right) \cdot z_k$. Due to Equation (3.5), this yields $z_k = \left(\frac{d-1}{d}\right)^{d-k-1}$. From Equations (3.4) and (3.2), we

get $w_k = \frac{1}{d-1} \cdot \left(\frac{d-1}{d}\right)^{d-k}$ and $y_k = 1 - \left(\frac{d-1}{d}\right)^{d-k}$, respectively. Finally, we verify that Equation (3.1) holds:

$$\begin{aligned} y_k - w_k + z_k &= \\ 1 - \left(\frac{d-1}{d}\right)^{d-k} - \frac{\left(\frac{d-1}{d}\right)^{d-k}}{d-1} + \left(\frac{d-1}{d}\right)^{d-k-1} &= \\ 1 + \left(\frac{d-1}{d}\right)^{d-k-1} \left(-\frac{d-1}{d} - \frac{1}{d} + 1\right) &= 1. \end{aligned}$$

Evaluating the goal function:

$$\begin{aligned} R_d + \sum_{k=1}^{d-1} y_k \cdot \frac{R_k}{d} &= \\ n \left(\left(\frac{d-1}{d}\right)^{d-1} + \frac{\sum_{k=1}^{d-1} \left(\frac{d-1}{d}\right)^{k-1}}{d} - \left(\frac{d-1}{d}\right)^d \right) &= \\ = n \left(1 - \left(\frac{d-1}{d}\right)^d \right). \end{aligned}$$

4 Online Capital Investment

In this section, we study the online Capital Investment problem. We must produce many units of a commodity at minimum cost, where orders for units arrive online. We have a set of machines, where each machine m_i has a capital cost c_i and production cost p_i . At any time, the algorithm can choose to buy any machine for cost c_i . The algorithm incurs a production cost of p_i if it uses machine m_i to produce one unit of the commodity. The goal is to minimize the total cost: the sum of capital costs plus production costs.

We give a randomized lower bound that is arbitrarily close to e for the Capital Investment problem. In addition, our techniques enable us to give a different lower bound for each input configuration (i.e., values for machine capital costs and production costs) that is tight against each such configuration.

THEOREM 4.1. *There is no randomized algorithm with a competitive ratio strictly better than e for the Capital Investment problem.*

Proof.

Sequence definition: The setting for the problem is a set of n machines where machine m_i has a capital cost of $i+1$ and a production cost of 2^{-i^2} . Our input sequence consists of n phases, where in phase k the online algorithm needs to produce a total of 2^{k^2} products. That is, we introduce $2^{k^2} - 2^{(k-1)^2}$ orders for units in phase k , with two orders being introduced in the first phase. Clearly, OPT is $k+2$ in phase k .

Primal variables: Let $x_{k,i}$ be the fraction bought of the i^{th} machine in the k^{th} phase. Let $q_{k,i}$ be the fraction of products produced by the i^{th} machine in the k^{th} phase. Let c be a variable representing the competitive ratio guarantee.

The primal linear program:

$$\begin{aligned} \min c & \\ \text{s.t.: } \sum_{r=1}^k x_{r,i} &\geq q_{k,i} && \text{cons. } y_{k,i} \\ \sum_{i=1}^n q_{k,i} &= 1 && \text{cons. } w_k \\ \sum_{r=1}^k \sum_{i=1}^n (i+1)x_{r,i} + \sum_{i=1}^n 2^{k^2-i^2} q_{k,i} &\leq c(k+2) && \text{cons. } z_k \\ c, x_{k,i}, q_{k,i} &\geq 0, && \end{aligned}$$

where the constraint $y_{k,i}$ is $\forall k, i \in [n]$, w_k is $\forall k \in [n]$, z_k is $\forall k \in [n]$, and the nonnegativity constraints are $\forall k, i \in [n]$. The constraint $y_{k,i}$ captures the fact that a machine cannot be used more than the amount paid for it. The constraint w_k implies that we need to use a machine to produce units. The constraint z_k captures the competitive ratio guarantee of the online algorithm. Namely, in phase k , the online algorithm's total capital costs plus production costs must not exceed $c \cdot OPT$. Note that we allow the online algorithm to produce all products in phase k , which allows the online algorithm to be refunded for the production done in previous phases. This of course does not decrease OPT and may only decrease the competitive ratio.

The dual linear program:

$$\begin{aligned} \max \sum_{k=1}^n w_k & \\ \text{s.t.: } \sum_{k=1}^n (k+2) \cdot z_k &\leq 1 && \text{cons. } c \\ (i+1) \sum_{r=k}^n z_r &\geq \sum_{r=k}^n y_{r,i} && \text{cons. } x_{k,i} \\ y_{k,i} &\geq w_k - z_k \cdot 2^{k^2-i^2} && \text{cons. } q_{k,i} \\ y_{k,i}, z_k &\geq 0, && \end{aligned}$$

where constraints depending on k or i are $\forall k, i \in [n]$. Note that we can omit constraint c by replacing the goal function with $\frac{\sum w_k}{\sum (k+2)z_k}$, which is achieved by normalizing all dual variables appropriately.

Intuition for the dual variables assignment:

A natural assignment is $y_{k,i} = w_k$ for $k \leq i$ and

0 otherwise (since $2^{k^2-i^2}$ should dominate w_k/z_k), along with $w_k = \frac{1}{k}$. This assignment yields ($\forall k \geq i : k, i \in [n]$)

$$(i+1) \sum_{r=k}^n z_r \geq \sum_{r=k}^i w_r \quad \text{cons. } x_{k,i}.$$

We view constraints $x_{k,i}$ as if they were ‘continuous’ (i.e., for large $i \geq k$). In doing so, we consider two differentiable functions $f(x)$ and $g(x)$, where $f'(k)$ and $g'(i)$ represent approximately the variables z_k and w_i , respectively. With this, we get $\forall i \geq k$

$$(i+1) \int_k^{n+1} f'(x) dx \geq \int_k^{i+1} g'(x) dx \iff$$

$$f(n+1) - f(k) \geq \frac{\ln\left(\frac{i+1}{k}\right)}{i+1},$$

which holds for

$$f(x) = -\frac{1}{e \cdot x}, \quad f'(x) = \frac{1}{e \cdot x^2},$$

since $\ln(x)/x \leq 1/e$ for $x \geq 1$. We assume that $f(n) \rightarrow 0$ as $n \rightarrow \infty$ and ignore it.

Formal feasible dual variables assignment:

Let ϵ be a small constant, the assignment is:

- $y_{k,i} = w_k$, for $k \leq i \leq n$ and 0 otherwise.
- $z_k = \frac{1}{k(k+1)}$, for all $k \leq n$.
- $w_k = e \cdot (1 - \epsilon) \ln\left(\frac{k+1}{k}\right)$, for $k \leq n \cdot \epsilon$ and 0 otherwise.

First, we verify that constraints $q_{k,i}$ hold for $k > i \geq 1$ (they trivially hold for $k \leq i$):

$$w_k - z_k \cdot 2^{k^2-i^2} \leq e \ln\left(\frac{k+1}{k}\right) - \frac{2^{2 \cdot k-1}}{k(k+1)} \leq 0 = y_{k,i}.$$

Now we verify that constraints $x_{k,i}$ hold for all $i \leq \epsilon \cdot n$ (since for $i > \epsilon \cdot n$ the left hand side of constraint $x_{k,i}$ increases while the right hand side remains the same as i increases). By assigning to the dual variables:

$$(i+1) \sum_{r=k}^n z_r - \sum_{r=k}^i w_r =$$

$$(i+1) \left(\frac{1}{k} - \frac{1}{n+1} \right) - e \cdot \ln\left(\frac{i+1}{k}\right) \cdot (1 - \epsilon) =$$

$$(1 - \epsilon) \left(\frac{i+1}{k} - e \cdot \ln\left(\frac{i+1}{k}\right) \right) +$$

$$\left(\frac{\epsilon \cdot (i+1)}{k} - \frac{i+1}{n+1} \right) \geq 0,$$

where the inequality holds since $x \geq e \cdot \ln(x)$ for all $x \geq 1$ and $k \leq i \leq \epsilon \cdot n$.

Evaluating the goal function: Recall that $H(n) = 1 + \frac{1}{2} + \dots + \frac{1}{n}$ denotes the n^{th} harmonic number. We get a lower bound of

$$\frac{\sum_{k=1}^n w_k}{\sum_{k=1}^n (k+2)z_k} = \frac{e \cdot \ln(n \cdot \epsilon) \cdot (1 - \epsilon)}{\sum_{k=1}^n \frac{k+2}{k(k+1)}} \geq$$

$$\frac{e \cdot \ln(n \cdot \epsilon) \cdot (1 - \epsilon)}{H(n) + C_1} \rightarrow e \cdot (1 - \epsilon),$$

where C_1 and ϵ are some constants. This gives the theorem.

5 Conclusions

We introduce a technique for proving online lower bounds using duality. Using our systematic method, we show how to construct new, tight lower bounds for three diverse online problems: Vector Bin Packing, Ad-auctions, and Capital Investment. We are also able to reconstruct many existing online lower bounds. We are certain that the techniques we develop here can have far-reaching implications, and can be used to improve existing lower bounds along with proving new, tight lower bounds as well.

References

- [1] Noga Alon, Baruch Awerbuch, Yossi Azar, Niv Buchbinder, and Joseph (Seffi) Naor. The online set cover problem. In *Proceedings of the 35th annual ACM Symposium on Theory of Computing*, 2003.
- [2] Nir Andelman and Yishay Mansour. Auctions with budget constraints. In *Proceedings of the 9th Scandinavian Workshop on Algorithm Theory*, 2004.
- [3] John Augustine, Sandy Irani, and Chaitanya Swamy. Optimal power-down strategies. *SIAM Journal on Computing*, 37(5):1499–1516, 2008.
- [4] Yossi Azar, Yair Bartal, Esteban Feuerstein, Amos Fiat, Stefano Leonardi, and Adi Rosén. On capital investment. In *Proceedings of the 23rd International Colloquium on Automata, Languages, and Programming*, 1996.
- [5] Yossi Azar, Umang Bhaskar, Lisa Fleischer, and Debmalaya Panigrahi. Online mixed packing and covering. In *Proceedings of the 24th annual ACM-SIAM Symposium on Discrete Algorithms*, 2013.
- [6] Yossi Azar, Ilan Reuven Cohen, Amos Fiat, and Alan Roytman. Packing small vectors. In *Proceedings of the 27th annual ACM-SIAM Symposium on Discrete Algorithms*, 2016.
- [7] Yossi Azar, Ilan Reuven Cohen, Seny Kamara, and Bruce Shepherd. Tight bounds for online vector bin

- packing. In *Proceedings of the 45th annual ACM Symposium on Theory of Computing*, 2013.
- [8] Nikhil Bansal, Niv Buchbinder, and Joseph (Seffi) Naor. A primal-dual randomized algorithm for weighted paging. In *Proceedings of the 48th annual IEEE Symposium on Foundations of Computer Science*, 2007.
- [9] Nikhil Bansal, Niv Buchbinder, and Joseph (Seffi) Naor. Metrical task systems and the k-server problem on hsts. In *Proceedings of the 37th International Colloquium on Automata, Languages, and Programming*, 2010.
- [10] Nikhil Bansal, Niv Buchbinder, and Joseph (Seffi) Naor. Towards the randomized k-server conjecture: A primal-dual approach. In *Proceedings of the 21st annual ACM-SIAM Symposium on Discrete Algorithms*, 2010.
- [11] Nikhil Bansal, Alberto Caprara, and Maxim Sviridenko. Improved approximation algorithms for multidimensional bin packing problems. In *Proceedings of the 47th annual IEEE Symposium on Foundations of Computer Science*, 2006.
- [12] Nikhil Bansal, Marek Eliáš, and Arindam Khan. Improved approximation for vector bin packing. In *Proceedings of the 27th annual ACM-SIAM Symposium on Discrete Algorithms*, 2016.
- [13] Reuven Bar-Yehuda and Shimon Even. A linear-time approximation algorithm for the weighted vertex cover problem. *Journal of Algorithms*, 2(2):198–203, 1981.
- [14] Yair Bartal, Francis Y. L. Chin, Marek Chrobak, Stanley P. Y. Young, Wojciech Jawor, Ron Lavi, Jiří Sgall, and Tomáš Tichý. Online competitive algorithms for maximizing weighted throughput of unit jobs. In *Proceedings of the 21st Symposium on Theoretical Aspects of Computer Science*, 2004.
- [15] Avrim Blum and Jason D. Hartline. Near-optimal online auctions. In *Proceedings of the 16th annual ACM-SIAM Symposium on Discrete Algorithms*, 2005.
- [16] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, 2005.
- [17] Niv Buchbinder, Kamal Jain, and Joseph (Seffi) Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *Proceedings of the 15th annual European Symposium on Algorithms*, 2007.
- [18] Niv Buchbinder and Joseph (Seffi) Naor. The design of competitive online algorithms via a primal-dual approach. *Foundations and Trends in Theoretical Computer Science*, 3(2–3):93–263, 2009.
- [19] Chandra Chekuri and Sanjeev Khanna. On multi-dimensional packing problems. In *Proceedings of the 10th annual ACM-SIAM Symposium on Discrete Algorithms*, 1999.
- [20] Edward G. Coffman Jr., János Csirik, Gábor Galambos, Silvano Martello, and Daniele Vigo. Bin packing approximation algorithms: survey and classification. In *Handbook of Combinatorial Optimization*, pages 455–531. Springer, 2013.
- [21] Peter Damaschke. Nearly optimal strategies for special cases of on-line capital investment. *Theoretical Computer Science*, 302(1):35–44, 2003.
- [22] Ran El-Yaniv, Amos Fiat, Richard Karp, and Gordon Turpin. Competitive analysis of financial games. In *Proceedings of the 33rd annual IEEE Symposium on Foundations of Computer Science*, 1992.
- [23] Ran El-Yaniv and Richard M. Karp. The mortgage problem. In *Proceedings of the 2nd Israeli Symposium on Theory of Computing and Systems*, 1993.
- [24] Gabor Galambos and Gerhard J. Woeginger. Online bin packing – a restricted survey. *Zeitschrift für Operations Research*, 42(1):25–45, 1995.
- [25] Minos N. Garofalakis and Yannis E. Ioannidis. Scheduling issues in multimedia query optimization. *ACM Computing Surveys*, 1995.
- [26] Minos N. Garofalakis and Yannis E. Ioannidis. Multi-dimensional resource scheduling for parallel queries. In *Proceedings of 1996 ACM SIGMOD International Conference on Management of Data*, 1996.
- [27] Michel X. Goemans and David P. Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, 24(2):296–317, 1995.
- [28] Kamal Jain, Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing lp. *Journal of the ACM*, 50(6):795–824, 2003.
- [29] Bala Kalyanasundaram and Kirk R. Pruhs. An optimal deterministic algorithm for online b-matching. *Theoretical Computer Science*, 233(1):319–325, 2000.
- [30] Anna R. Karlin, Mark S. Manasse, Lyle A. McGeoch, and Susan Owicki. Competitive randomized algorithms for nonuniform problems. *Algorithmica*, 11(6):542–571, 1994.
- [31] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd annual ACM Symposium on Theory of Computing*, 1990.
- [32] Amir Levi and Boaz Patt-Shamir. Non-additive two-option ski rental. *Theoretical Computer Science*, 584:42–52, 2015.
- [33] Zvi Lotker, Boaz Patt-Shamir, and Dror Rawitz. Rent, lease, or buy: Randomized algorithms for multislope ski rental. *SIAM Journal on Discrete Mathematics*, 26(2):718–736, 2012.
- [34] Mohammad Mahdian and Amin Saberi. Multi-unit auctions with unknown supply. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, 2006.
- [35] Mohammad Mahdian and Qiqi Yan. Online bipartite matching with random arrivals: An approach

based on strongly factor-revealing lps. In *Proceedings of the 43rd annual ACM Symposium on Theory of Computing*, 2011.

- [36] Aranyak Mehta and Debmalya Panigrahi. Online matching with stochastic rewards. In *Proceedings of the 53rd annual IEEE Symposium on Foundations of Computer Science*, 2012.
- [37] Aranyak Mehta, Amin Saberi, Umesh V. Vazirani, and Vijay V. Vazirani. Adwords and generalized on-line matching. In *Proceedings of the 46th annual IEEE Symposium on Foundations of Computer Science*, 2005.
- [38] Adam Meyerson, Alan Roytman, and Brian Tagiku. Online multidimensional load balancing. In *International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*. Springer Berlin Heidelberg, 2013.
- [39] Rina Panigrahy, Kunal Talwar, Lincoln Uyeda, and Udi Wieder. Heuristics for vector bin packing. *Microsoft Research T.R.*, 2011.
- [40] Meikel Poess and Raghunath Othayoth Nambiar. Energy cost, the key challenge of today’s data centers: a power consumption analysis of tpc-c results. *Proceedings of the VLDB Endowment*, 2008.
- [41] Prabhakar Raghavan. A statistical adversary for on-line algorithms. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 7:79–83, 1992.
- [42] Sara Robinson. Computer scientists optimize innovative ad auction. *SIAM news*, 38(3):3, 2005.
- [43] André van Vliet. An improved lower bound for on-line bin packing algorithms. *Information Processing Letters*, 43(5):277–284, 1992.
- [44] Yajun Wang and Sam Chiu-wai Wong. Two-sided online bipartite matching and vertex cover: Beating the greedy algorithm. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming*, 2015.
- [45] Guiqing Zhang, Chung Keung Poon, and Yin-feng Xu. The ski-rental problem with multiple discount options. *Information Processing Letters*, 111(18):903–906, 2011.

A Suboptimal Assignment for Vector Bin Packing

Intuition for the dual variables assignment: We note that by assigning $y_i = 1$ for all i leads to the following non-optimal solution (with a value approaching 2). To find an assignment for variables $z_{k,j}$, we look at the constraints corresponding to $x_{i,j}$ as if they were ‘continuous’ (i.e., for large $i \geq j$). With this interpretation, we get:

$$i \cdot f'_j(i) + \int_j^i f'_j(x) dx \geq 1 \iff x \cdot f'_j(x) + f_j(x) - f_j(j) \geq 1.$$

By solving this differential equation (assuming equality) with the boundary condition $f_j(j) = 0$, we get:

$$f_j(x) = 1 - \frac{j}{x}, \quad f'_j(x) = \frac{j}{x^2}.$$

Formal feasible dual variables assignment:

$$y_i = 1, \quad z_{k,j} = \begin{cases} \frac{j-1}{k \cdot (k-1)} & \text{if } k > 1, k \geq j, \\ 1 & \text{if } k = j = 1, \\ 0 & \text{otherwise.} \end{cases}$$

With this assignment, we need to verify that constraint $x_{i,j}$ is feasible for all $i \geq j$. For $j = 1$ and for $i \geq j$, the constraint corresponding to $x_{i,1}$ easily holds, since both sides of the inequality evaluate to 1. For $j > 1$, we get:

$$i \cdot \frac{j-1}{i \cdot (i-1)} + \sum_{r=j}^{i-1} \frac{j-1}{r \cdot (r-1)} = \frac{j-1}{i-1} + (j-1) \cdot \left(\frac{1}{j-1} - \frac{1}{i-1} \right) = 1.$$

Evaluating the goal function: Note that $\sum_{k=j}^d z_{k,j} = 1 - \frac{j-1}{d}$, and hence the value of the goal function is

$$\frac{\sum_{i=1}^d y_i}{\sum_{k=1}^d \sum_{j=1}^d z_{k,j}} = \frac{d}{d - \sum_{j=1}^d \frac{j-1}{d}} = \frac{d}{d - \frac{d-1}{2}} \rightarrow 2.$$

B Simulating Randomized Algorithms with Deterministic Fractional Algorithms

In this section, given a randomized algorithm, we show how to simulate it using a deterministic fractional algorithm. In particular, we guarantee that the performance of the simulated deterministic fractional algorithm for any sequence is the same as the expected performance of the randomized algorithm. Therefore, a lower bound for any deterministic fractional algorithm is also a lower bound for any randomized algorithm. In general, such a simulation is done by considering the randomized algorithm’s expected behavior, and viewing it as a deterministic fractional algorithm where the fractions are the expected probabilities of the randomized algorithm (with the appropriate definition of a fractional algorithm).

For the online Ad-auctions problem and the online Capital Investment problem, this is straightforward. In particular, for the online Ad-auctions problem, the expectation of the randomized algorithm’s

behavior is equivalent to splitting each item. For the online Capital Investment problem, the randomized algorithm's behavior is equivalent to purchasing fractional machines.

For the Vector Bin Packing problem, we need to be more careful about our definition of a fractional algorithm, since the behavior of the randomized algorithm includes both splitting vectors randomly and opening bins randomly. A randomized algorithm that splits vectors randomly can be viewed as a deterministic algorithm splitting vectors (according to the expectation), while opening bins randomly can be viewed as opening a fractional amount of bins (according to the expectation).