How to Allocate Goods in an Online Market?

Yossi Azar[1], Niv Buchbinder[2], and Kamal Jain[3]

[1] Tel-Aviv University, Tel-Aviv, 69978, Israel. *
[2] Microsoft Research, New England.
[3] Microsoft Research, Redmond.

Abstract. We study an online version of Fisher's linear case market. In this market there are m buyers and a set of n dividable goods to be allocated to the buyers. The utility that buyer i derives from good j is u_{ij}. Given an allocation $\hat{U}$ in which buyer i has utility $\hat{U}_i$ we suggest a quality measure that is based on taking an average of the ratios $U_i/\hat{U}_i$ with respect to any other allocation U. We motivate this quality measure, and show that market equilibrium is the optimal solution with respect to this measure. Our setting is online and so the allocation of each good should be done without any knowledge of the upcoming goods.
We design an online algorithm for the problem that is only worse by a logarithmic factor than any other solution with respect to our proposed quality measure, and in particular competes with the market equilibrium allocation. We prove a tight lower bound which shows that our algorithm is optimal up to constants. Our algorithm uses a primal dual convex programming scheme. To the best of our knowledge this is the first time that such a scheme is used in the online framework.
We also discuss an application of the framework in display advertising business in the last section.

1 Introduction

Allocating goods to buyers in a way that maximizes social welfare and fairness guarantees has been studied extensively. One well established framework for achieving a desirable allocation is the market equilibrium concept [1]. In a general market setting known as Fisher's linear case market we are given m buyers and n divisible goods. Each buyer has a budget e_i. The utility functions are linear, which means that buyer i derives a utility u_{ij} out of good j. It is well known that in this market (as well as many other more general markets) there exist prices for the goods and a corresponding (equilibrium) allocation of the goods to the buyers with several desirable properties. First, all goods are fully allocated and all buyers fully extract their budgets. Second, buyers are only assigned goods that belong to their optimal basket. That is, buyers are only allocated goods that maximize the utility per price for the buyer. This allocation has two additional nice properties. First, when considering splitting a buyer's budget into two buyers whose utilities are the same as the original buyer, the sum of allocations of the market equilibrium solution after the split is also optimal before the split. This property shows the robustness of the market equilibrium allocation since buyers do not benefit

* Research supported in part by the Israel Science Foundation.

from simulating themselves using several "smaller" buyers. A second property is that scaling all utilities of some buyer by a constant does not change the allocation, which is a desirable property because utilities of buyers are incomparable. This again means that buyers cannot benefit from boosting up their utilities. As early results capture only existence of the allocation, computing the market equilibrium allocation (and prices) for various markets in polynomial time received much attention recently [13, 11, 15, 18, 12].

Consider an allocation $\hat{U}$ in which buyer i has utility $\hat{U}_i$. We would like to measure the quality of the allocation with respect to any other allocation U. Since utilities of different buyers are incomparable and may be scaled up or down, it is only meaningful to look at the ratios $U_i/\hat{U}_i$, for $1 \leq i \leq m$. A natural way to measure the quality of $\hat{U}$ with respect to some other allocation U is to select some average, for example an arithmetic average, and to look at the average of $U_i/\hat{U}_i$. The lower the average is the better $\hat{U}$ is. When considering the quality and fairness of $\hat{U}$, this evaluation can be done with respect to any other allocation U. It is then natural to evaluate the quality of the allocation $\hat{U}$ as the maximum over all possible allocations U of these averages. Our proposed quality measure for $\hat{U}$ is therefore the following:

$$\max_{U} \left\{ \operatorname*{avg}_{i=1}^{m} \frac{U_i}{\hat{U}_i} \right\} \tag{1}$$

where avg can be any weighted average with the budgets, e_i, of the buyers playing the role of weights.

Looking for the best or fairest allocation then corresponds to looking for an allocation that minimizes the maximum. That is, we are looking for U^* that is the solution to the following:

$$U^* = \arg\min_{\hat{U}} \left\{ \max_{U} \left\{ \operatorname*{avg}_{i=1}^{m} \frac{U_i}{\hat{U}_i} \right\} \right\}$$

We explore this definition for several natural averages: arithmetic, geometric and harmonic averages. A simple observation shows that market equilibrium is the optimal offline allocation for both the harmonic and the geometric averages. This follows directly from the convex program suggested by Eisenberg and Gale [14] as a way to compute market equilibrium in Fisher's linear case. Interestingly, we show that market equilibrium is also optimal for the **arithmetic** average. See Section 2.1 for a detailed proof. These results mean that the market equilibrium allocation achieves a value of 1 (which is optimal) with respect to quality measure (1) which motivates further our definition.

Market equilibrium is, indeed, a desirable allocation. However, in many settings the allocation of goods should be done in an online fashion, where goods are arriving one-by-one, and previous allocation decisions are irrevocable.

As one motivating example one may think of a wireless router (base station) with many users. At any point in time there is a quality of transmission for each user, and the router should decide how to split the bandwidth between the users [20]. The problem is, of course, online, where every time slot corresponds to a new product. Achieving a fair allocation in this online environment is a natural goal. A second example is allocating impressions to advertisers on the internet, so as to maximize social welfare and fairness. It is worth mentioning that our setting considers fractional allocations while impressions are indivisible goods. However, fractional allocations (like the one that is generated later by our algorithm) can be simulated in many cases using simple randomization with only additional small loss. See further discussion in Section 5. The online setting raises a natural question of whether it is possible to achieve some of the desirable qualities of the offline market equilibrium allocation within an online framework.

1.1 Our results

We study the quality of allocations that may be achieved in an online setting with respect to our proposed quality measure. Our main contribution is designing an online allocation mechanism with the following properties:

Theorem 1. *Let U be any allocation of the goods to the buyers. Our online algorithm computes an allocation $\hat{x}_{ij}$ with utilities $\hat{U}_i$ such that:*

$$\max_U \left\{ \sum_{i=1}^{m} e_i \cdot \frac{U_i}{\hat{U}_i} \right\} \leq 1 + \ln m + \ln n + \ln \frac{u_{\max}}{u_{\min}}$$

where $\frac{u_{\max}}{u_{\min}}$ is the ratio of the maximum utility a buyer derives from a good divided by the minimum **non zero** *utility the buyer derives from any other good.*

Remark 11 *While Theorem 1 refers to an arithmetic average, we immediately get the same performance guarantee with respect to the harmonic and geometric averages by the Arithmetic-Geometric-Harmonic Means Inequality.*

This result shows that even in an online setting we may achieve an allocation whose average is only worse by a logarithmic factor than any offline allocation. We show that the performance of our algorithm is tight up to constants by proving the following lower bounds on the performance of any online algorithm even when all budgets are equal. We prove two lower bounds. A tight lower bound for the arithmetic and geometric averages, and an almost tight lower bound for the harmonic average.

Theorem 2. *Let $\hat{U}_i$ be the utilities achieved by any online algorithm and let U_i^* be the utilities of the market equilibrium allocation, then there exists an instance with $n \geq m^2$ such that:*

- $\frac{1}{m} \sum_{i=1}^{m} \frac{U_i^*}{\hat{U}_i} \geq \left(\prod_{i=1}^{m} \frac{U_i^*}{\hat{U}_i} \right)^{1/m} = \Omega \left(\min \left\{ m, \ln n + \ln \frac{u_{\max}}{u_{\min}} \right\} \right)$
- $\frac{1}{\frac{1}{m} \sum_{i=1}^{m} \frac{\hat{U}_i}{U_i^*}} = \Omega \left(\min \left\{ \frac{m}{\ln m}, \frac{\ln n + \ln \frac{u_{\max}}{u_{\min}}}{\ln \ln n + \ln \ln \frac{u_{\max}}{u_{\min}}} \right\} \right)$

where $\frac{u_{\max}}{u_{\min}}$ is the ratio of the maximum utility a buyer derives from a good divided by the minimum **non zero** *utility the buyer derives from any other good.*

Techniques: Whenever a good arrives, our algorithm computes an allocation to the newly arrived good, by solving a small linear program that can be interpreted as invoking Karush, Kuhn, Tucker (KKT) optimality conditions of the convex program with respect to the current allocation and the utility function of the new good. Along with the allocation of the good, we get a dual variable that is later used in our analysis. To the best of our knowledge this is the first time that such a primal dual convex programming scheme is used in the online framework.

Another main difference from previous works is that while the performance measure in most works is essentially a ratio of sums, our performance measure is a sum (average) of ratios. We believe that this measure of quality is very reasonable and is applicable to other scenarios as well.

1.2 Previous results

Existence of a market equilibrium allocation in a very general setting was proved by Arrow and Debreu [1]. Algorithmic aspects of the offline linear case of Fisher's model [6, 19] were studied in [13]. They designed a polynomial time algorithm that computes prices for the goods and the market equilibrium allocation. This was done by designing an algorithm that solves the convex program suggested by Eisenberg and Gale [14] (see our preliminaries). Computing a market equilibrium allocation in the offline case for other markets has also been studied recently [15] (See also [18]). Allocation of goods in an online fashion was studied in many different settings [4, 16, 17, 9]. For instance, Blum et al. consider a setting in which sellers and buyers are trading a single commodity. Sell/Buy bids arrive online with an expiration time, and the goal of the auctioneer is to match these bids so as to maximize revenue, or social welfare.

Competing against offline solutions in an adversarial setting was studied in many settings [5]. The closest to our setting is the problem of scheduling jobs on unrelated machines. In this problem, there is a set of m machines and a set of n jobs. The load of job j on machine i is l_{ij}. The basic problem is to minimize the load on the most loaded machine. Awerbuch et al. [2] designed an online $O(\log m)$-competitive algorithm for this problem. This result is tight. Later an equivalent primal-dual interpretation of the algorithm was shown in [7, 8] (See also [10]). A more general measure of performance of minimizing the l_p norm for any p was studied in [3]. They showed that the greedy algorithm is $O(p)$-competitive which is the best possible. Notice, however, several important differences between the load balancing problem and our problem. First, the current problem is a maximization problem and not a minimization problem. Second, while the performance measure in all these works is essentially a ratio of sums, our performance measure is a sum (average) of ratios.

2 Preliminaries

We study here Fisher's linear case model. Our market consists of a set of m buyers, and n divisible goods. Let e_i be the budget of buyer i. The budget represents here the importance of the buyer, and for simplicity of representation we normalize e_i so that $\sum_{i=1}^{m} e_i = 1$. The utilities functions of the goods to each buyer are linear. Let u_{ij} be the utility buyer i derives from good j. Let $u_{\max,i} \triangleq \max_{j=1}^{n}\{u_{ij}\}$ be the maximum utility the buyer gets from a good. Let $u_{\min,i} \triangleq \min_{j|u_{ij}>0}\{u_{ij}\}$ be the minimum non-zero utility from a good of buyer i. Let $\frac{u_{\max}}{u_{\min}} \triangleq \max_{i=1}^{m}\left\{\frac{u_{\max,i}}{u_{\min,i}}\right\}$ be the maximum ratio of utilities over the buyers. Without loss of generality we assume here that for each good j there exists a buyer i such that $u_{ij} > 0$ (otherwise the good may be discarded). Also, for each buyer i there exists a good j such that $u_{ij} > 0$ (otherwise buyer i gets utility 0 in any allocation). In an allocation, let x_{ij} be the amount of good j allocated to buyer i. Given an allocation the utility derived by buyer i is $U_i \triangleq \sum_{j=1}^{n} u_{ij}x_{ij}$. We study here an online setting in which goods arrive one-by-one in an online fashion. Upon arrival of a good the algorithm should decide how to allocate it to the buyers, and this decision is irrevocable.

The market equilibrium allocation: It is well known that the following convex program suggested first by Eisenberg and Gale [14] can be used in order to compute the market equilibrium allocation in Fisher's linear model:

$$\max \sum_{i=1}^{m} e_i \cdot \ln(U_i)$$

Subject to:

$$\begin{aligned} \forall 1 \le i \le m \quad & U_i = \textstyle\sum_{j=1}^{n} u_{ij}x_{ij} \\ \forall 1 \le j \le n \quad & \textstyle\sum_{i=1}^{m} x_{ij} \le 1 \\ \forall 1 \le i \le m, 1 \le j \le n \quad & x_{ij} \ge 0 \end{aligned}$$

Let x^*_{ij} be the optimal solution to the convex program. We may define a set of lagrangian variables p_j for each good j in the program. These variables may be interpreted as prices for the goods. The KKT optimality conditions define a relationship between the optimal values of x^*_{ij} and p_j. First, all p_j are strictly positive (assuming that each good has an interested buyer). In the optimal allocation x^*_{ij} each good is fully allocated and also the two following conditions are satisfied:

Optimality conditions:

1. For each buyer i and item j: $\frac{p_j}{e_i} \ge \frac{u_{ij}}{\sum_{j=1}^{n} u_{ij}x^*_{ij}}$
2. For each buyer i and item j: $x^*_{ij} > 0 \Rightarrow \frac{p_j}{e_i} = \frac{u_{ij}}{\sum_{j=1}^{n} u_{ij}x^*_{ij}}$

Using these conditions it is also possible to prove that $\sum_{j=1}^{n} p_j = \sum_{i=1}^{m} e_i = 1$ which means that the prices p_j are clearing the market (extract all the budgets of the buyers).

2.1 Our performance measure

As explained in the introduction we chose our performance measure for an allocation $\hat{U}$ to be:

$$\max_{U}\left\{\operatorname*{avg}_{i=1}^{m} \frac{U_i}{\hat{U}_i}\right\} \tag{2}$$

We choose to concentrate on studying the arithmetic, geometric and harmonic averages. We prove first that the market equilibrium allocation is optimal with respect to all these measures, which gives motivation for our study of this performance measure.

Lemma 1. *Let x^*_{ij} be the market equilibrium allocation and let U^*_i be the utilities of the buyers. Then for any other allocation x_{ij} with utilities U_i to the buyers:*

$$\frac{1}{\sum_{i=1}^m e_i \cdot \frac{U_i^*}{U_i}} \le \prod_{i=1}^m \left(\frac{U_i}{U_i^*}\right)^{e_i} \le \sum_{i=1}^m e_i \cdot \frac{U_i}{U_i^*} \le 1.$$

We also note that the maximum ratio of 1 *is tight since it is tight for the allocation* $U_i = U_i^*$.

Proof. We remark first that it is clear that:

$$\prod_{i=1}^m \left(\frac{U_i}{U_i^*}\right)^{e_i} \le 1$$

since the convex programming above actually maximizes $(\prod_{i=1}^m U_i^{e_i})$. Therefore, it is left to prove only the last inequality. Let p_j be the market clearance prices as computed by the convex program.

$$\sum_{i=1}^m e_i \cdot \frac{U_i}{U_i^*} = \sum_{i=1}^m e_i \cdot \sum_{j=1}^n \frac{u_{ij}x_{ij}}{\sum_{j=1}^n u_{ij}x^*_{ij}} \le \sum_{i=1}^m e_i \cdot \sum_{j=1}^n \frac{p_j}{e_i} x_{ij} \tag{3}$$

$$= \sum_{j=1}^n p_j \sum_{i=1}^m x_{ij} = \sum_{j=1}^n p_j = 1 \tag{4}$$

where Inequality (3) follows by the first optimality condition and the last equality follows by the second fact that the prices clear the market.

3 The Algorithm

In this section we design and analyze our main algorithm for finding an allocation of the goods in an online fashion. Note the similarity between the way our algorithm computes the current allocation and the optimality conditions of the convex linear program used to computed the market equilibrium allocation. Our algorithm works as follows.

When a new good j arrives with utilities u_{ij}:

- For any $k < j$ let $\hat{x}_{ik}$ be the allocation of the algorithm to the previous goods.
- Compute an allocation for the current good by solving the following optimization problem:
  $\min p_j$

  $\sum_{i=1}^m x_{ij} \le 1$

  For each buyer i such that $u_{ij} > 0$:
  $\frac{p_j}{e_i} \ge \frac{u_{ij}}{u_{ij}x_{ij} + \sum_{k<j} u_{ij}\hat{x}_{ij}}$

We remark that for all $k < j$, $\hat{x}_{ik}$, that is the allocation of the online algorithm of the previous goods, are constants. The only variables in the program are x_{ij} of the current good and p_j. We state the algorithm in this certain way that is convenient for our proof. However, we illustrate here a different way of viewing the algorithm that shows an easy way to compute the allocation and gives an additional intuition. To find the allocation one may define a variable $\tilde{p}_j \triangleq \frac{1}{p_j}$. Minimizing p_j corresponds to maximizing $\tilde{p}_j$. Using this variable we get the following simple linear program.

$\max \tilde{p}_j$

$\sum_{i=1}^{m} x_{ij} \leq 1$

For each buyer i and good j such that $u_{ij} > 0$:
$\tilde{p}_j \leq \frac{x_{ij}}{e_i} + \frac{\sum_{k<j} u_{ij}\hat{x}_{ij}}{u_{ij} \cdot e_i}$

In the constraints imposed on each buyer and good the second term is simply a constant which is the utility obtained so far by the user normalized by the utility of the current good and the budget (importance) of the buyer. It is possible to obtain an optimal solution to the LP by increasing x_{ij} in rate proportional to e_i to all buyers for which the RHS of the constraint is minimized. This is done until $\sum_{i=1}^{m} x_{ij} = 1$. Viewing the process that way we obtain a simple observation that we later use in our proof.

Observation 31 *Let j' be the first good for which $u_{ij'} > 0$, then $\hat{x}_{ij'} \geq e_i$.*

More intuition for our algorithm can be obtained by considering the simpler special case where the utilities are either 0 and 1, and all budgets are equal. For this special case the algorithm reduces to a "water level" algorithm that tries to balance the utilities of the buyers. Therefore, the algorithm can be viewed as an adaptation of the water level algorithm to this more complex setting.

Additional properties of the algorithm: The main idea of the algorithm is to invoke the offline KKT optimality conditions with respect to our current solution. This allow us later to bound the quality of our solution as a function of the dual variables obtained in the process. We assume here without loss of generality that for any good j there is a buyer i for which $u_{ij} > 0$. This means that for every good j, p_j that is computed by our algorithm is strictly larger than 0. Given this condition the following important lemma states an immediate property of the algorithm with respect to the $\hat{x}_{ij}$ and the dual values p_j it computes. The lemma is an online version of the complementary slackness conditions that are obtained offline.

Lemma 2. *Let $\hat{x}_{ij}$ and p_j be the allocation and the values p_j computed during the execution of the online algorithm then:*

1. *For each buyer i and good j: $\frac{p_j}{e_i} \geq \frac{u_{ij}}{\sum_{k=1}^{j} u_{ik}\hat{x}_{ik}}$, / or $u_{ij} = \sum_{k=1}^{j} u_{ik}\hat{x}_{ik} = 0$.*
2. *For each buyer i and good j: $\hat{x}_{ij} > 0 \Rightarrow \frac{p_j}{e_i} = \frac{u_{ij}}{\sum_{k=1}^{j} u_{ik}\hat{x}_{ik}}$*

3.1 Analysis of the algorithm

In this section we prove our main result bounding the quality achieved by the algorithm. We prove that the online algorithm achieves an allocation that is at most logarithmic factor worse than any other allocation with respect to quality measure 1 and in particular from the best allocation with respect to this measure which is the market equilibrium allocation.

Theorem 3. *Let U be any allocation of the goods to the buyers. Our online algorithm computes an allocation $\hat{x}_{ij}$ with utilities $\hat{U}_i$ such that:*

$$\max_U \left\{ \sum_{i=1}^m e_i \cdot \frac{U_i}{\hat{U}_i} \right\} \leq 1 + \ln m + \ln n + \ln \frac{u_{\max}}{u_{\min}}$$

where $\frac{u_{\max}}{u_{\min}}$ is the ratio of the maximum utility a buyer derives from a good divided by the minimum **non zero** *utility the buyer derives from any other good.*

Proof. Let $p_1, p_2, \ldots, p_n$ be the dual variables computed by the online algorithm. Let $\hat{x}_{ij}$ be the amount of good j that was allocated by our algorithm to buyer i. Let x_{ij} be the amount of item j allocated to buyer i in any other fixed solution. We assume here without loss of generality that $x_{ij} > 0$ only when $u_{ij} > 0$. By the properties of our algorithm we get that.

$$\begin{aligned}
\sum_{i=1}^m e_i \cdot \frac{U_i}{\hat{U}_i} &= \sum_{i=1}^m e_i \cdot \frac{\sum_{j=1}^n u_{ij} x_{ij}}{\sum_{k=1}^n u_{ik} \hat{x}_{ik}} = \sum_{i=1}^m e_i \cdot \sum_{j=1}^n \frac{u_{ij} x_{ij}}{\sum_{k=1}^n u_{ik} \hat{x}_{ik}} \\
&\leq \sum_{i=1}^m e_i \cdot \sum_{j=1}^n \frac{u_{ij} x_{ij}}{\sum_{k=1}^j u_{ik} \hat{x}_{ik}} && (5) \\
&\leq \sum_{i=1}^m e_i \cdot \sum_{j=1}^n \frac{p_j}{e_i} x_{ij} = \sum_{j=1}^n p_j \sum_{i=1}^m x_{ij} && (6) \\
&\leq \sum_{j=1}^n p_j && (7) \\
&= \sum_{j=1}^n \sum_{i=1}^m p_j \hat{x}_{ij} = \sum_{i=1}^m e_i \cdot \sum_{j=1}^n \frac{p_j}{e_i} \hat{x}_{ij} && (8) \\
&= \sum_{i=1}^m e_i \cdot \sum_{j=1}^n \frac{u_{ij} \hat{x}_{ij}}{\sum_{k=1}^j u_{ik} \hat{x}_{ik}} && (9)
\end{aligned}$$

Inequality (5) follows by summing over fewer values. Inequality (6) follows by the first property in Lemma 2. Inequality (7) follows since in any allocation $\sum_{i=1}^n x_{ij} \leq 1$. Equality (8) follows since in our allocation $\sum_{i=1}^n \hat{x}_{ij} = 1$. Finally, Equality (9) follows by the second property of Lemma 2. To bound Inequality (9) we prove the following claim.

Claim. For each buyer $1 \le i \le m$:

$$\sum_{j=1}^{n} \frac{u_{ij}\hat{x}_{ij}}{\sum_{k=1}^{j} u_{ik}\hat{x}_{ik}} \le 1 + \ln\left(\sum_{j=1}^{n} u_{ij}\hat{x}_{ij}\right) - \ln\left(u_{ij'}\hat{x}_{ij'}\right)$$
$$= 1 + \ln(\hat{U}_i) - \ln\left(u_{ij'}\hat{x}_{ij'}\right)$$

where j' is the first good for which $\hat{x}_{ij'} > 0$.

Proof. The proof is by induction on the number goods n. Let j' be the first good for which $\hat{x}_{ij'} > 0$. For this good the LHS as well as the RHS are 1 and so the claim holds. It is easy to see that both sides of the inequality change only when $\hat{x}_{ij} > 0$, thus, we only consider such goods. Assume that the claim holds for $\ell - 1$ goods for which $\hat{x}_{ij} > 0$. Then For the ℓth good we get that:

$$\sum_{j=1}^{\ell} \frac{u_{ij}\hat{x}_{ij}}{\sum_{k=1}^{j} u_{ik}\hat{x}_{ik}} \le 1 + \ln\left(\sum_{j=1}^{\ell-1} u_{ij}\hat{x}_{ij}\right) - \ln\left(u_{ij'}\hat{x}_{ij'}\right) + \frac{u_{i\ell}\hat{x}_{i\ell}}{\sum_{k=1}^{\ell} u_{ik}\hat{x}_{ik}} \tag{10}$$
$$\le 1 + \ln\left(\sum_{j=1}^{\ell} u_{ij}\hat{x}_{ij}\right) - \ln\left(u_{ij'}\hat{x}_{ij'}\right) \tag{11}$$

Inequality (10) follows by the induction hypothesis. Inequality 11 reduces to proving that:

$$\ln\left(\sum_{j=1}^{\ell-1} u_{ij}\hat{x}_{ij}\right) - \ln\left(\sum_{j=1}^{\ell} u_{ij}\hat{x}_{ij}\right) + \frac{u_{i\ell}\hat{x}_{i\ell}}{\sum_{k=1}^{\ell} u_{ik}\hat{x}_{ik}}$$
$$= \ln\left(1 - \frac{u_{i\ell}\hat{x}_{i\ell}}{\sum_{k=1}^{\ell} u_{ik}\hat{x}_{ik}}\right) + \frac{u_{i\ell}\hat{x}_{i\ell}}{\sum_{k=1}^{\ell} u_{ik}\hat{x}_{ik}} \le 0$$

The final inequality is true since for any $0 < a < 1$, $\ln(1-a) \le -a$.

Plugging Claim 3.1 into Inequality (9), we get:

$$\sum_{i=1}^{m} e_i \cdot \frac{U_i}{\hat{U}_i} \le \sum_{i=1}^{m} e_i \cdot \left(1 + \ln \hat{U}_i - \ln\left(u_{ij'}\hat{x}_{ij'}\right)\right)$$
$$= 1 + \sum_{i=1}^{m} e_i \cdot \left(\ln \hat{U}_i - \ln\left(u_{ij'}\hat{x}_{ij'}\right)\right) \tag{12}$$

To bound (12) we bound the second term. By observation 31 the worst case is when for i and j' $u_{ij'} = u_{\min,i}$ where $u_{\min,i}$ is the minimum non zero utility of buyer i, and all other buyers received no utility so far. In this case the algorithm assigns a fraction of e_i of the good to buyer i. Thus, $-\ln\left(u_{ij'}\hat{x}_{ij'}\right) \le \ln\frac{1}{e_i} - \ln u_{min,i}$. and we get:

$$
\begin{aligned}
\sum_{i=1}^{m} e_i \cdot \frac{U_i}{\hat{U}_i} &\leq 1 + \sum_{i=1}^{m} e_i \cdot \left(\ln \frac{1}{e_i} + \ln \frac{\hat{U}_i}{u_{\min,i}} \right) \\
&\leq 1 + \sum_{i=1}^{m} e_i \ln \frac{1}{e_i} + \ln n + \ln \frac{u_{\max}}{u_{\min}} \qquad (13) \\
&\leq 1 + \ln m + \ln n + \ln \frac{u_{\max}}{u_{\min}}. \qquad (14)
\end{aligned}
$$

Inequality (13) follows since $\hat{U}_i \leq n u_{\max,i}$, where $u_{\max,i}$ is the maximum utility of buyer i. Inequality (14) follows since the entropy on m different values is at most $\log m$. This concludes our proof.

4 Lower Bound

In this section we show that the performance of our algorithm is almost tight. We show a lower bound that is tight up to constants on the performance of any online allocation algorithm for the geometric and arithmetic averages. For the harmonic average we show an almost tight lower bound. Due to lack of space the proof of Theorem 4 appears in Appendix A.

Theorem 4. *Let $\hat{U}_i$ be the utilities achieved by any online algorithm and let U_i^* be the utilities of the market equilibrium allocation, then there exists an instance with $n \geq m^2$ such that:*

- $\frac{1}{m} \sum_{i=1}^{m} \frac{U_i^*}{\hat{U}_i} \geq \left(\prod_{i=1}^{m} \frac{U_i^*}{\hat{U}_i} \right)^{1/m} = \Omega \left(\min \left\{ m, \ln n + \ln \frac{u_{\max}}{u_{\min}} \right\} \right)$
- $\frac{1}{\frac{1}{m} \sum_{i=1}^{m} \frac{\hat{U}_i}{U_i^*}} = \Omega \left(\min \left\{ \frac{m}{\ln m}, \frac{\ln n + \ln \frac{u_{\max}}{u_{\min}}}{\ln \ln n + \ln \ln \frac{u_{\max}}{u_{\min}}} \right\} \right)$

where $\frac{u_{\max}}{u_{\min}}$ is the ratio of the maximum utility a buyer derives from a good divided by the minimum **non zero** *utility the buyer derives from any other good.*

5 Further Discussion

We believe that an interesting part of our work is the suggestion of our quality measure. This quality measure which is essentially an average of ratios is quite natural in the setting of Fisher's linear case market, and we believe it may be beneficial in other scenarios as well. The choice of what average to use is quite flexible and may be explored further in terms of fairness. While we consider an online setting, it is also interesting to study simple offline dynamics that lead to good allocations in terms of our proposed quality measure, as an example consider the following applications.

Consider the following setting of display advertising. Suppose there is a company which sells display advertising and has a very detailed technology of ad-targeting.

Whenever a user views any of its webpages, the company can make reliable prediction of the worth of the user to the advertisers.

Since the advertisers are relatively less sophisticated than the company, the company offers a limited bid expressive language for the advertisers to describe their ad-targeting needs. At the time of placing an advertising order, the advertisers could describe the age, gender, geographical location etc, of the desired audience being targeted. During the run time, i.e. when users are visiting the websites, the company is likely to have more detailed information about the users, and using the sophisticated machine learning algorithms could make prediction about the worth of an user to the advertisers.

The advertising order in the display industry is usually guaranteed, i.e., once the order is accepted then the company would have to serve the order or pay severe penalties of not fulfilling the orders. An advertising order usually has a minimum impression count which must be fulfilled, and a maximum dollar budget which could be charged. Usually when the company accepts an advertising order, it has sure that with high probability the order can be fulfilled. Usually the company has access to low quality ad inventory which could also be used to fulfill the orders.

Note that the company can't keep allocating the low quality ad inventory to an advertiser. Doing so may mean that the advertiser may not return back to buy more advertising in the future. So the company needs some kind of method to fairly allocate the advertising space.

In summary, the company has N advertisers with their budgets described. Users arrive in an online fashion. Whenever a user arrives, the company predicts the value of the user (i.e., utility) to the advertisers. Based on this predicted value, the company assigns the user to one of the advertisers. Since a large number of users are expected to arrive, one could assume fractional allocation too, which could be converted into integral allocation by using randomized rounding on a large number of users.

Note that, unlike the AdWords problem, the entire budget will be charged here, since the obligation towards an advertiser is assumed be matched with low quality ad inventory, if needed. The AdWords problem makes advertisers satisfied by not charging a part of their budget when the ad-inventory is not delivered. In display advertising there is no such flexibility, and the advertisers must be made happy by allocating proportionately a high quality ad-inventory. This is precisely the model studied in this paper.

Acknowledgement

We would like to thank Gagan Goel for suggesting the wireless router motivation, and also acknowledge to attend a talk presented by Nitish Karula. A discussion among the audience members with the speaker inspired the application mentioned in the discussion section.

References

1. K. Arrow and G. Debreu. Existence of an equilibrium for competitive economy. *Econometrica*, pages 265–290, 1954.
2. J. Aspnes, Y. Azar, A. Fiat, S. A. Plotkin, and O. Waarts. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *J. ACM*, 44(3):486–504, 1997.

3. Baruch Awerbuch, Yossi Azar, Edward F. Grove, Ming-Yang Kao, P. Krishnan, and Jeffrey Scott Vitter. Load balancing in the lp norm. In *Proc. of 36th FOCS*, pages 383–391, 1995.
4. Avrim Blum, Tuomas Sandholm, and Martin Zinkevich. Online algorithms for market clearing. *J. ACM*, 53(5):845–879, 2006.
5. Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998.
6. W. C. Brainard and H.E. Scarf. How to compute equilibrium prices in 1891. In *Cowles Foundations Discussion paper*, page 1270, 2000.
7. N. Buchbinder and J. Naor. Online primal-dual algorithms for covering and packing problems. In *13th Annual European Symposium on Algorithms*, 2005.
8. N. Buchbinder and J. Naor. Improved bounds for online routing and packing via a primal-dual approach. In *Proceedings of the 47th annual ieee Symposium on Foundations of Computer Science (FOCS 2006)*, 2006.
9. Niv Buchbinder, Kamal Jain, and Joseph (Seffi) Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *ESA '07: Proceedings of the 15th Annual European Symposium*, pages 253–264, 2007.
10. Niv Buchbinder and Joseph (Seffi) Naor. The design of competitive online algorithms via a primaldual approach. *Foundations and Trends in Theoretical Computer Science*, 3:93–263, 2009.
11. Deeparnab Chakrabarty, Nikhil R. Devanur, and Vijay V. Vazirani. New results on rationality and strongly polynomial time solvability in eisenberg-gale markets. In *WINE*, pages 239–250, 2006.
12. Nikhil R. Devanur and Ravi Kannan. Market equilibria in polynomial time for fixed number of goods or agents. In *FOCS*, pages 45–53, 2008.
13. Nikhil R. Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V. Vazirani. Market equilibrium via a primal-dual-type algorithm. In *FOCS*, pages 389–395, 2002.
14. E. Eisenberg and D. Gale. Consensus of subjective probabilioties: The pari-mutuel method. *Annual of mathematical statistics*, 30:165–168, 1959.
15. Kamal Jain and Vijay V. Vazirani. Eisenberg-gale markets: algorithms and structural properties. In *STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 364–373, 2007.
16. Mohammad Mahdian and Amin Saberi. Multi-unit auctions with unknown supply. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 243–249, 2006.
17. Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5):22, 2007.
18. Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
19. H. Scarf. The computation of economic equilibria (with collaboration of t. hansen). In *Cowles foundation monograph No. 24*, 1973.
20. Alexander L. Stolyar. Greedy primal-dual algorithm for dynamic resource allocation in complex networks. *Queueing Syst. Theory Appl.*, 54(3):203–220, 2006.

A Proof of Theorem 4

Proof. We first consider the case of utilities 0 or 1. We are going to consider a family of adversarial setting parameterized for a some number of buyers m and number of goods $m \leq n \leq 2^m$ (the number of goods eventually will not be exactly n). For each value n we are going to partition the set of m buyers into disjoint sets of $\log n$ buyers each.

The set of buyers in each disjoint set will have positive utility to a disjoint set of goods. Therefore, both the optimal market equilibrium allocation and any online allocation algorithm will allocate all the goods within a set only to the buyers in this set. We next describe the adversarial setting within a single set of $\log n$ buyers.

For any such subset of $\log n$ buyers there are $\log n$ rounds in which goods arrive. In round 1 a single good arrives and all $\log n$ buyers have utility 1 for the good. Since only one unit of the good can be allocated there must exist a buyer to which only at most $1/\log n$ unit of the good was allocated. Without loss of generality assume this buyer is buyer 1. Therefore, at the second round 2 goods arrive and buyers 2 to $\log n$ have utility 1 for the two goods. In general in round $1 \leq i \leq \log n$, 2^{i-1} goods arrive and the subset of buyers that have utility 1 for the goods is the subset of $\log n - i + 1$ buyers to which the online algorithm allocated the least fraction of goods so far. Observe that the total number of goods that arrive in all rounds is at most $1 + 2 + \ldots, +2^{\log n - 1} \leq n$ (in all subsets the number of goods is more than n, but this will not change our results, see later in the proof).

Next, we analyze the performance of any online allocation algorithm in this adversarial setting. The utility that the online algorithm allocated to the ith "poorest" buyer in the subset is no more than:

$$\hat{U}_i \leq \frac{1}{\log n} + \frac{2}{\log n - 1} + \ldots + \frac{2^{i-1}}{\log n - i + 1} \leq \frac{2^i}{\log n - i + 1}$$

However, in the market equilibrium solution the best solution that maximizes the geometric average of utilities is to allocate the first good in full to the first buyer, the next two goods in full to the second buyer and so on. In general the utility of the ith "poorest" buyer in the market equilibrium allocation is:

$$U_i^* = 2^{i-1}$$

Therefore, we get that:

$$\frac{U_i^*}{\hat{U}_i} \geq \frac{\log n - i + 1}{2}$$

Considering the first claim of the lower bound inside each subset:

$$\prod_{i=1}^{\log n} \frac{U_i^*}{\hat{U}_i} \geq \frac{(\log n)!}{n} \geq \frac{1}{n} \left(\frac{\log n}{3}\right)^{\log n} \tag{15}$$

where the inequality is for large enough n. Therefore, considering all $m/\log n$ subsets, and considering the fact that the real number of goods in all subsets is n times the number of subsets, that is $m/\log n$ which doesn't our logarithmic lower bound, we get that:

$$\left(\prod_{i=1}^{m} \frac{U_i^*}{\hat{U}_i}\right)^{1/m} \geq \left(\frac{1}{n}\left(\frac{\log n}{3}\right)^{\log n}\right)^{\frac{m}{\log n} \cdot \frac{1}{m}} \geq \frac{\log n}{6} = \Omega(\log n)$$

If $n \geq 2^m$ then the bound when $n = 2^m$ is $\Omega(m)$ and we are done.

Next, we consider the second claim in the lower bound. For each subset:

$$\sum_{i=1}^{\log n} \frac{\hat{U}_i}{U_i^*} \leq \sum_{i=1}^{\log n} \frac{2}{\log n - i + 1} \leq 2(1 + \ln \log n) \tag{16}$$

Therefore, summing over all subsets we get:

$$\frac{1}{\frac{1}{m} \sum_{i=1}^{m} \frac{\hat{U}_i}{U_i^*}} \geq \frac{1}{\frac{1}{m} \cdot \frac{m}{\log n} 2\left(1 + \ln \log n\right)} = \Omega\left(\frac{\log n}{\log \log n}\right)$$

If $n \geq 2^m$ then the bound when $n = 2^m$ is $\Omega(m/\log m)$ and we are done.

To see the bound for general utility values, we may replace each 2^i items by a single item and have the utility of each buyer 2^i. This reduces the total number of goods so that $n = m$, while the bound remains the same.