Efficient Allocation of Free Stuff

Yossi Azar
Tel Aviv University
azar@tau.ac.il

Allan Borodin
University of Toronto
bor@cs.toronto.edu

Michal Feldman
Tel Aviv University
michal.feldman@cs.tau.ac.il

Amos Fiat
Tel Aviv University
fiat@tau.ac.il

Kineret Segal
Tel Aviv University
kineretsegal@mail.tau.ac.il

July 31, 2018

Her yigitin bir yoqurt yeyisi vardir
(Free vinegar is sweeter than honey)

Turkish Proverb

Abstract

We study online matching settings with selfish agents when everything is free. Inconsiderate agents break ties arbitrarily amongst equal maximal value available choices, even if the maximal value is equal to zero. Even for the simplest case of zero/one valuations, where agents arrive online in an arbitrary order, and agents are restricted to taking at most one item, the resulting social welfare may be negligible for a deterministic algorithm. This may be surprising when contrasted with the $1/2$ approximation of the greedy algorithm, analogous to this setting, except that agents are considerate (i.e., they don't take zero-valued items). We overcome this challenge by introducing a new class of algorithms, which we refer to as prioritization algorithms. We show that upgrading a random subset of the agents to "business class" already improves the approximation to a constant. For more general valuations, we achieve a constant approximation using $\log n$ priority classes, when the valuations are known in advance. We extend these results to settings where agents have additive valuations and are restricted to taking up to some $q \geq 1$ items. Our results are tight up to a constant.

1 Introduction

Almost universally, market efficiency is achieved by setting prices on goods. We consider an online market setting where buyers arrive over time. Mechanisms for social welfare in such markets have been studied in, e.g., [14, 12, 17]. A measure of the quality of such mechanisms is "how well do they approximate the maximal social welfare?" As in these previous studies, we assume that the order of arrival is adversarial.

The key issue considered in this paper is *"What efficiency can be achieved in online markets where goods are given away for free?"*. Unfortunately, it is easy to see that without prices only a negligible fraction of the social welfare is achievable. This holds even if we restrict agents to take at most one item (or some q items). Moreover, this poor efficiency holds even if the agent valuations are zero/one, given that the agents are *inconsiderate*. An inconsiderate agent will choose to take an item of no value to them if they have no better option. Also, an inconsiderate agent will break ties amongst equally valuable items in an arbitrary (and inconsiderate) manner. There is much research to suggest agents may in fact behave so. E.g., see [20, 15] and references therein.

The key idea in this paper is to categorize agents into priority classes, where agents from a higher priority class always precede those from a lower class, but the order within a class is arbitrary.

We begin by considering the simplest setting with zero/one valuations and inconsiderate agents about whom we know nothing.

Free distribution to Inconsiderate Strangers

Consider the following scenario: prior to departing on vacation, we seek to distribute our remaining food to passers-by (agents). Every agent is unit demand with zero/one valuations, but we know nothing about their preferences, nor do we know the order in which they arrive. Every agent, upon arrival, is allowed to choose a single item from those remaining.

If agents are "well behaved" and only choose an item that they like then the resulting distribution is a maximal matching (in a bipartite unweighted graph of agents and items, where an edge indicates that the agent likes the item), which is know to be a 1/2 approximation to the maximum matching.

However, human nature being what it is [20, 15], agents who see nothing of value to themselves may be *inconsiderate* and may take an item for which they have no perceived value [1]

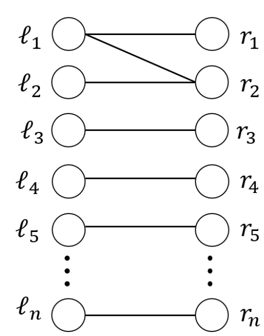

Figure 1: An example showing that inconsiderate agents may result in low social welfare.

It is easy to construct an example where such inconsiderate agents "steal" items of value from subsequent passers-by, despite having no value for these items themselves (albeit, such subsequent agents need know nothing of this). To see this, consider the scenario depicted in Figure 1. There are n agents, $\ell_1, \ldots, \ell_n$, and n items, $r_1, \ldots, r_n$, and every agent ℓ_i has value 1 for item r_i, with the exception of agent ℓ_1 that has value 1 also for item r_2. Suppose agents arrive in an increasing order of their indices $(\ell_1, \ldots, \ell_n)$, and agent ℓ_1 arbitrarily chooses item r_2 over r_1, thereafter every agent ℓ_i, $i < n$, takes item r_{i+1}, and ℓ_n takes r_1. The resulting social welfare is negligible when compared with the maximum social welfare (1 instead of n).

[1] An alternate explanation for agents taking "zero value" items (and for breaking ties arbitrarily) is that the true valuations are a slight perturbation of some underlying ground truth, see, e.g.,[21].

A natural approach to overcome this problem is to reject (or delay) such problematic agents that have no item of value remaining. If we knew the items that agents care about, we could prioritize agents with an item of value remaining, delaying others and still get a maximal matching.

Note that the problem presented above assumes no prior knowledge about the agents, and, moreover, as no prices are used, strategic agents may claim to like everything when in fact they like nothing from the leftover items . Thus, it seems on first glance that blindly prioritizing some agents over others is useless.

We argue that this intuition is both true and false: We show that prioritizing some agents deterministically gives negligible social welfare (See Section 3). In contrast, if there is a perfect matching, then by selecting a random set of prioritized agents, we obtain a $1/4$ approximation to the optimum. Moreover, if there is an assignment of items to agents where αn agents get an item (for some $\alpha \leq 1$) — then, by selecting a random set of prioritized agents — we obtain an $\alpha/4$ approximation to the optimum (see Theorem 1). Unfortunately, this is also tight (up to a constant). For large α this is fine but is not great if α is small. Moreover, this problem is inherent for more general valuations.

We remark that this trivial prioritization algorithm (prioritize a random subset of agents) is oblivious in the sense that it knows nothing about the agents (i.e., the graph is unknown), the order of arrival, agent identities, how agents break ties, and what items are leftover. We also note that this prioritization algorithm can be run on the fly, where all agents arrive in some adversarial order and are classified into priority classes on the fly.

To give good approximations in the case of small α and for more general valuations we turn to a model of "Inconsiderate Friends". The distinction between strangers and friends is that for strangers we know nothing about their valuation for items whereas for friends we know how much they like each item (but not how they break ties). This allows us to get much better approximations than in the case of strangers.

In particular, knowing agent valuations gives a constant factor approximation for zero/one valuations and for arbitrary α (see Theorem 3). The more interesting case is that of general unit demand valuations (the agent valuation for every item is arbitrary).

Free distribution to Inconsiderate Friends with Unit Demand Valuations

The rules of the game are that we can prioritize our friends (with the goal of maximizing social welfare). For example, we can invite some set of friends in the morning and another in the evening. The morning friends arrive in some arbitrary unknown order and arbitrarily break ties, the same holds for friends that arrive in the evening (and choose only amongst the morning leftovers). The key issue is the number of such priority classes. Obviously, having more classes yields better approximation ratios.

Here, agents are unit demand and valuations are described as an edge weighted bipartite graph. Our main result is that, for arbitrary such valuations, and assuming worst case order and worst case tie breaking, one can achieve a $c \cdot r/\log n$ approximation if allowed r priority classes, for some constant c, and that this is tight. See Theorem 7.

We remark that rather than insisting on unit demand valuations, our results also hold where agent valuations are completely arbitrary (even complementarities are allowed). In this case we still insist that an agent can take no more than one item. Critically, in this general valuation setting, the benchmark is not the maximal social welfare for the original valuations but rather the maximum social welfare achievable given that agents are restricted to taking one item. E.g., the value for a single shoe which could be much less than the value of a pair of shoes.

Furthermore, we consider additional extensions such as additive valuations in which agents have additive valuations, and they are restricted to take up to $q \geq 1$ items rather than only one item. (Inconsiderate agents will always take the full allotment). See Section 5.

1.1 Related Work

1.1.1 Online bipartite matching

The unweighted online matching problem can be represented by a bipartite graph, where nodes on the left represent agents, nodes on the right represent items, and the existence of an edge between an agent and an item means that the agent has value 1 for the item. In such problems, agents arrive over time, choosing an arbitrary adjacent remaining item. It is well known that irrespective of how agents make their choices, this process results in a maximal matching, thus yields at least half of the maximum matching. This is the best deterministic algorithm for unweighted graphs [13].

In their seminal paper, Karp, Vazirani and Vazirani [13] show that a randomized algorithm performs better. In particular, by imposing a random preference order on the items, the greedy algorithm gives at least $1 - 1/e$ of the optimal matching.

If agents have arbitrary valuations for items and can choose only one item (represented as a weighted bipartite matching), in general no guarantees on the efficiency can be obtained. Special cases have been considered in the literature [8, 1]. More generally, online bipartite matching has been an active area of theoretical computer science research for almost 30 years. The survey by Mehta [16] provides a excellent overview as to the various variants of online bipartite matching with applications to online advertising.

1.1.2 Posted pricing for known valuations

In the context of posted pricing, one should distinguish between considerate and inconsiderate tie breaking. If ties are broken appropriately, then Walrasian pricing exists for all gross substitute valuations [11]. This means that all items are assigned prices, and agents arrive sequentially, each offered a specific most desired bundle. Given such considerate tie breaking, such a process results in maximum welfare.

In [9] this approach has been generalized for arbitrary valuations, yielding half of the optimal welfare. However, prices are now attached to bundles of items, rather than to individual items.

To deal with inconsiderate tie breaking, Cohen-Addad *et al.* [6] give a dynamic variant of Walrasian pricing for unit-demand valuations, that achieves optimal welfare. With static posted prices, one can achieve half of the optimal welfare, but no more than 2/3.

1.1.3 Posted pricing for Bayesian settings

Feldman *et al.* [10] show that if the valuations are drawn (independently) from known probability distributions over submodular valuations, then half of the optimal welfare can be obtained in expectation using posted pricing. This work was later extended to more general stochastic settings, using the framework of prophet inequality [7].

1.1.4 Mechanisms without Money

The question of maximizing social welfare *without recourse to prices* has previously been studied in numerous settings such as facility location [19], common goods [3], cake cutting [18], social choice functions [5], and kidney exchange [2].

1.1.5 Relation to Priority Model

The priority model [4] was introduced to model greedy or more generally myopic algorithms. In the fixed order priority model, every input is given a distinct priority. Our *prioritization model* allows for a finer grained approach distinguishing intermediate problems between the standard online model and the priority model. The parameter of interest is the number of priority classes.

2 Model and Preliminaries

We model agent valuations using an edge weighted bipartite graph $G = (L, R; E)$, where $R = \{r_1, \ldots, r_m\}$ represents the set of items, and $L = \{\ell_1, \ldots, \ell_n\}$ the set of agents. The weight $w(e)$ of an edge $e = (\ell_i, r_j)$ from $\ell_i \in L$ to $r_j \in R$ is the value agent ℓ_i has for item r_j. We sometime abuse notation and write $w(i, j)$ to denote the weight of the edge (ℓ_i, r_j).

In this paper, items never have prices, everything is free. However, agents are restricted in how many items they can take. We first consider allowing one item, and discuss generalizations subsequently.

We consider *prioritization* algorithms where agents can be assigned to some priority class. Agents with higher priority make their selection before agents of lower priority. The highest priority class is C_1, for multiple priority classes, agents belonging to priority class C_i choose items before agents belonging to priority classes C_j, $j > i$. Items that have been selected by some agent disappear and are unavailable for an agent to arrive subsequently.

Agents assigned to no priority class (the *plebeians*) are last to choose. Within a single priority class, (and within the plebeian class) the order of arrival and how ties are broken are determined adversarially.

For randomized algorithms, we consider an *oblivious* adversary. In our setting, this means that the adversary determines both a global order of arrival, and how agents break ties (should they arise). The adversary determines these issues, in advance, without knowing the random bits used by the algorithm. The global order determines how priority classes are ordered. The relative order of two agents that belong to the same priority class is implied by their relative position in the global order. (Likewise for plebeians).

Positive results (a lower bound on the social welfare), using at least one priority class, can ignore the plebeians in the analysis. Thus, for all positive results we simply ignore the plebeians. For negative results the contribution of the plebeians can be viewed as having one additional priority class.

We consider two different scenarios:

1. Strangers. In this setting nothing is known about the agents, the prioritization algorithm assigns [indistinguishable] agents to one of r priority classes, $C_1, \ldots, C_r$, or to none (the plebeians). *I.e.*, the agent/item graph is unknown, the order within a priority class is unknown, and how ties are broken is unknown. We say that such an algorithm is oblivious since it makes its decisions blindly, all agents are indistinguishable.

2. Friends. In this setting agent valuations to items are known before assigning agents to priority classes. *I.e.*, the agent/item graph is known, but not the order within a priority class nor how ties are broken. Prioritization algorithms assigns agents to one of r priority classes, $C_1, \ldots, C_r$, or to none (the plebeians).

3 Results for Inconsiderate Strangers

We consider a setting where we select some subset of agents to have priority, who can choose whatever item they want. Subsequently, the remaining (non-prioritized) agents can be allowed to choose items too[2]. Amongst the prioritized agents, the order in which they choose items is arbitrary.

As all strangers are indistinguishable, the only question is if to prioritize an agent (and allow the stranger to choose an item immediately) or not. As agents are asked no questions (and agents cannot be trusted anyway), and as agents are free to choose whatever maximizes their utility — it follows that strategic agents have no impact on the procedure and this process is inherently truthful.

It is trivial to observe that any algorithm that deterministically chooses what agents are prioritized results in an unbounded approximation ratio. Consider some deterministic priority algorithm, two agents a and b, and only item, for which one has value 1 and the other zero. As nothing is known, the prioritization algorithm will prioritize one of a and b, both of which are indistinguishable. Clearly, the agent chosen will have value zero for the item, yet, annoyingly, will choose it nonetheless.

Next, we consider a randomized prioritization algorithm (with a single priority class, in addition to the plebeians), and show the following:

Theorem 1. *For any $0 \leq \alpha \leq 1$, prioritizing every agent with probability $\frac{\alpha}{2}$ gives an $\frac{\alpha}{4}$ approximation to the size of the maximum matching, for any unweighted graph with a maximum matching of size $\geq \alpha n$.*

Proof. Given an unweighted graph G, fix some maximum matching M. Index the agents by the adversary global order starting with agent 1. Let $M(k)$, $1 \leq k \leq n$, be the item matched to agent k in the matching M, and $M(k) = \perp$ if no item was matched to agent k in M. Note that we do not know the adversary global order, nor do we know the $M(k)$'s.

For the purpose of analysis imagine that all agents arrive in the adversary global order where the ith event is the arrival of agent i. Agents are either allowed to make a choice or assigned to the plebeian class (and thus delayed). This assignment to the plebeian class is done on the fly.

We say that item j is unavailable after event i, $0 \leq i \leq n$ if it was taken by a prioritized agent j such that $j \in \{1, \ldots, i\}$.

For $0 \leq i < k \leq n$ define n_k^i as follows:

$$n_k^i := \begin{cases} 1 & \textit{if } M(k) \neq \perp \text{ and } M(k) \text{ is not available after event } i \\ 0 & \textit{otherwise.} \end{cases}$$

[2] In the analysis of the positive results (a lower bound on the social welfare) we assume that they give zero contribution to the social welfare, ergo, choosing not to prioritize an agent is equivalent to discarding the agent.

Let $S_i := \sum_{k>i} n_k^i$. The following holds:

$$S_{i+1} \leq S_i - n_{i+1}^i + I_{i+1} \tag{1}$$

where $I_{i+1} = 1$ if agent $i+1$ is prioritized and otherwise $I_{i+1} = 0$.

The difference $S_{i+1} - S_i$ consists of several components. $S_{i+1} - S_i$ clearly decreases by n_{i+1}^i, and if $I_{i+1} = 1$ this difference may increase by one — this happens when agent $i+1$ takes an item in $\{M(i+2), \ldots, M(n)\}$.

Taking expectation over (1), using linearity of expectation, and noting that agent $i+1$ takes an item (any item) with probability $\frac{\alpha}{2}$, we get that:

$$\mathbb{E}[S_{i+1}] \leq \mathbb{E}[S_i] - \mathbb{E}[n_{i+1}^i] + \frac{\alpha}{2},$$

or equivalently

$$\mathbb{E}[S_{i+1}] - \mathbb{E}[S_i] \leq -\mathbb{E}[n_{i+1}^i] + \frac{\alpha}{2}.$$

Taking the sum of i from 0 to $n-1$, we get that

$$\mathbb{E}[S_n] - \mathbb{E}[S_0] \leq -\sum_{i=0}^{n-1} \mathbb{E}[n_{i+1}^i] + n \cdot \frac{\alpha}{2}.$$

Note that $\mathbb{E}[S_n] = \mathbb{E}[S_0] = 0$ and hence

$$\sum_{i=0}^{n-1} \mathbb{E}[n_{i+1}^i] \leq n \cdot \frac{\alpha}{2}. \tag{2}$$

Let R_i be the size of the matching after event i, the sequence R_i is [weakly] ascending. Define $J_i = 1$ if $M(i) = \perp$ and zero otherwise. We now show that

$$R_{i+1} \geq R_i + I_{i+1} \cdot \big(1 - n_{i+1}^i - J_{i+1}\big), \tag{3}$$

by the following case analysis

- If $I_{i+1} = 0$ or $n_{i+1}^i = 1$ or $J_{i+1} = 1$ then (3) follows directly from monotonicity of R_i.
- The only remaining case is when $I_{i+1} = 1$ and both $n_{i+1}^i = 0$ ($M(i+1)$ was available after event i) and $J_{i+1} = 0$ ($M(i+1) \neq \perp$), and then the size of the matching increases by one.

Taking the expectation over (3) we get that

$$\mathbb{E}[R_{i+1}] \geq \mathbb{E}\Big[R_i + I_{i+1} \cdot \big(1 - n_{i+1}^i - J_{i+1}\big)\Big].$$

It follows from linearity of expectation and the fact that I_{i+1} is independent of n_{i+1}^i that

$$\mathbb{E}[R_{i+1}] \geq \mathbb{E}[R_i] + \mathbb{E}[I_{i+1}]\Big(1 - \mathbb{E}[n_{i+1}^i] - J_{i+1}\Big),$$

or equivalently,

$$\mathbb{E}[R_{i+1}] - \mathbb{E}[R_i] \geq \frac{\alpha}{2}\Big(1 - \mathbb{E}[n^i_{i+1}] - J_{i+1}\Big).$$

Taking the sum for i from 0 to $n-1$, we get that

$$\mathbb{E}[R_n] - \mathbb{E}[R_0] \geq n \cdot \frac{\alpha}{2} - \frac{\alpha}{2}\sum_{i=0}^{n-1}\mathbb{E}[n^i_{i+1}] - \frac{\alpha}{2} \cdot (1-\alpha)n,$$

note that $\mathbb{E}[R_0] = 0$. Using the bound for $\mathbb{E}[n^i_{i+1}]$ from (2) we have that

$$\mathbb{E}[R_n] \geq \frac{n\alpha}{2} - \frac{\alpha}{2} \cdot \frac{n\alpha}{2} - \frac{n\alpha}{2} + \frac{n\alpha^2}{2} = \frac{n\alpha^2}{4}. \tag{4}$$

Now, using (4) we give a bound on the approximation ratio of the algorithm:

$$\frac{\text{Alg}}{\text{Opt}} \geq \frac{n\alpha^2/4}{n\alpha} = \frac{\alpha}{4}.$$

where Opt is the size of the maximum weighted matching and Alg is the expected size of the matching achieved by the algorithm. □

Remark: Even if α is unknown, one can use standard techniques to guess the value of α to within some constant factor and lose a factor of $\log n$ on the competitive ratio.

We now show that the approximation ratio given in Theorem 1 is tight up to a constant factor.

Theorem 2. *For an unweighted graph with a maximum matching of size αn, no prioritization algorithm can achieve an approximation to the maximum matching greater than α.*

Proof. Consider an instance with n agents and n items as depicted in Figure 2. Suppose agents break ties (amongst zero valued items) from top to bottom. The first αn agents to arrive choose items $r_1, ..., r_{\alpha n}$ so there is no point in allowing more than αn agents to take an item. The adversary can schedule αn random slots for agents $\ell_1, ..., \ell_{\alpha n}$. The best way to choose a priority class, in this case, is choosing a random subset of αn agents, resulting in an expected social welfare of value $\alpha^2 n$. Hence, the approximation ratio is α. □

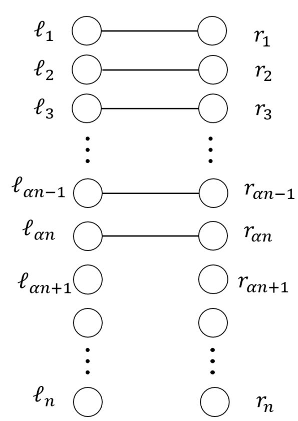

Figure 2: For an unweighted graph with a maximum matching of size αn, the approximation ratio is at least α. In this example agents break ties in favor of higher items (i.e., items with low indices).

We next turn to the problem of assigning agents, *whose valuations are known*, to [a small number of] priority classes. Clearly, it cannot be harder to prioritize agents if their valuations are known than if not. For zero/one valuations, this "friends" model (known valuations) improves the approximation above to a constant, if the size of the maximum matching in an unweighted graph is small (small α). Furthermore, this allows us to give good approximations to the value of the maximum matchings in the more general setting where agents have arbitrary valuations (not restricted to zero/one values).

4 Results for Inconsiderate Friends

Theorem 3. *For unweighted graphs (valuations zero/one), it is possible to choose a [random] subset of the agents as a higher priority class, and achieve a $\frac{1}{4}$ approximation to the size of the maximum matching, independent of the size of the matching (i.e., independent of α).*

Proof. First, we compute a maximum matching and exclude all unmatched agents, effectively this means that all remaining agents have a match. Then, from the remaining agents, prioritize each one with probability $\frac{1}{2}$. As all agents have a match, we can apply Theorem 1 with $\alpha = 1$ that yields a 1/4 approximation. □

This is almost tight:

Lemma 4. *For unweighted graphs (valuations zero/one), no prioritization algorithm with one priority classes can attain approximation ratio greater than $\frac{2}{3}$.*

Proof. Consider an instance with n agents and n items, as depicted in Figure 3. Agents are divided into $n/3$ sets $\{L_i\}_i$ of size 3, $L_i = \{\ell_i^1, \ell_i^2, \ell_i^3\}$ for $i = 1, \ldots, n/3$. Items are likewise divided into $n/3$ sets $\{R_i\}_i$ of size 3, $R_i = \{r_i^1, r_i^2, r_i^3\}$ for $i = 1, \ldots, n/3$. The set of edges is $E = \{(\ell_i^j, r_i^k) |\ k \leq j\}$.

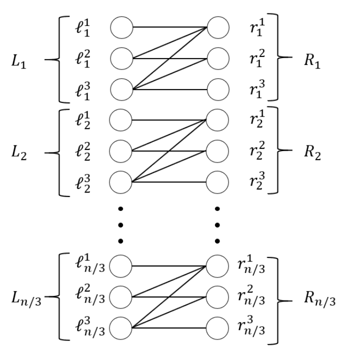

Figure 3: For an unweighted known graph, no prioritization algorithm with one priority class can achieve an approximation better than 2/3.

In the optimal solution, agent ℓ_i^j takes item r_i^j, hence each L_i gives a value of 3, and in total the value of the maximum matching is n. We claim that for any algorithm with a single priority class, there is an adversary global order and tie breaking rule such that the resulting value of each L_i is at most 2, hence in total the social welfare is $2n/3$, which gives an approximation ratio of 2/3.

The following global ordering of the agents and tie breaking rule implies the claim for every subset of agents: agents are ordered from high indices to low indices (bottom to top in figure), and always break ties in favor of items with lower indices (top items in figure). Let C be the priority class. For any set L_i, there are two options:

- if $L_i \cap C \neq \{\ell_i^1\}$, then agent ℓ_i^1 takes a zero valued item. Hence the value contributed by the agents in L_i is not greater than 2.
- if $L_i \cap C = \{\ell_i^1\}$, then ℓ_i^2 cannot take r_i^2 because by the time she arrives r_i^2 is already taken (by either ℓ_i^3 or other agent) and again the value contributed by the agents in L_i is not greater than 2.

□

We now show that for arbitrary valuations, appropriately upgrading some agents to business class gets an $\Omega(1/\log n)$ fraction of the social welfare.

Lemma 5. *For arbitrary valuations (described as a weighted graph), there exists a prioritization algorithm with only one priority class (in addition to the plebeians) that gives approximation ratio of $\Omega(\frac{1}{\log n})$.*

Proof. Fix a maximum weighted matching M, and let W denote the weight of M. Index the agents by adversary global order starting with agent 1. Let $M(k)$, $1 \leq k \leq n$, be the item matched to agent k in the matching M, and $M(k) = \perp$ if no item was matched to agent k in M. For any agent k such that $M(k) \neq \perp$, we denote the value agent k has for item $M(k)$ by $w(M(k)) = w(k, M(k))$.

Discard all agents with $w(M(k)) < W/(2n)$. Since these agents contribute in total at most $W/2$ to the value of the maximum matching, discarding these agents can decrease the value of the maximum matching by a factor of 2 at most. Assign agent k for which $M(k) \neq \perp$ and $w(M(k)) > W/(2n)$ to classes as follows: agent k belongs to class C_i, $0 \leq i \leq \log n + 1$, if and only if $W/2^{i+1} \leq w(M(k)) < W/2^i$ (where class C_0 includes also agent k s.t. $M(k) = W$, if exists). Let $C_{\max}$ denote the class with the highest contribution to the social welfare. Agents in $C_{\max}$ contribute at least $\frac{1}{2\log n}$ fraction of the value of the maximum matching.

Now, we prioritize only agents from $C_{\max}$: every such agent is prioritized with probability $\frac{1}{2}$. Let $C_{\max} = C_j$ for some $0 \leq j \leq \log n + 1$. Now, consider a thought experiment where edges (ℓ_k, r_m), $\ell_k \in C_{\max}$ have weight zero if $w(\ell_k, r_m) < W/2^{j+1}$ and weight $W/2^{j+1}$ if $w(\ell_k, r_m) \geq W/2^{j+1}$. The value of a matching on a subset of these "thought experiment" agents is no greater than the value of a matching on the same subset with the original values.

We can treat the input as if it was zero/one values (where $W/2^{j+1}$ plays the role of one), and prioritize agents as if they had zero/one values; this loses at most a factor 2 due to rounding. Applying Theorem 3, we lose another factor of 4 of the total contribution of $C_{\max}$. Thus, this gives a social welfare of at least $\frac{W}{16\log n}$. □

This result is asymptotically tight:

Theorem 6. *For arbitrary valuations, adding a business class (in addition to the default class) does not give an approximation ratio greater than $O(\frac{1}{\log n})$. Moreover, no algorithm that uses r priority classes can achieve an approximation ratio greater than $O(\frac{r}{\log n})$.*

Proof. Consider the bipartite graph depicted in Figure 4, with n agents, $\ell_1, \ldots, \ell_n$, and n items, $r_1, \ldots, r_n$. The set of edges is $E = \{(\ell_i, r_j) | \, 1 \leq j \leq i \leq n\}$, and every edge of the form (ℓ_i, r_j) has weight $\frac{1}{i}$. Clearly, in the maximum matching agent ℓ_i is matched to item r_i, resulting in a total weight of $H_n = \sum_{i=1}^{n} \frac{1}{i} \approx \ln n$.

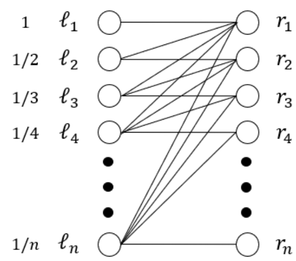

Figure 4: Agents on the left and items on the right. The number that appears to the left of an agent represents the corresponding weight of all the edges adjacent to this agent.

Fix an arbitrary subset of agents. We describe an adversary global order and tie breaking rule, and claim that that with these order and tie breaking rule the contribution of any subset to the social welfare is at most 1. This implies the theorem, since for a single priority class (i.e., partitioning the agents to two subsets) we get at most $2/\ln n$ of the optimal welfare, and with r priority classes we get at most $(r+1)/\ln n$ of the optimal welfare.

The global ordering of the agents and tie breaking rule are as follows: agents are ordered from high indices to low indices (bottom to top in figure), and always break ties in favor of items with lower indices (top items in figure). Consider a subset of agents that forms a priority class $C = \{i_1, \ldots, i_k\}$, where $i_j < i_{j+1}$ for $j = 1, \ldots, k-1$ Clearly, there is some index, i_{k^*}, such that for all $j \geq k^*$ agent i_j takes an item of positive value, and for all $j < k^*$ agent i_j takes a zero valued item.

Assume without loss of generality that $i_{k^*}, \ldots, i_k$ are consecutive agents (i.e., for all $k^* \leq j \leq k-1$, $i_j + 1 = i_{j+1}$). This is without loss since agents with lower indices contribute higher values. Therefore, $i_k - i_{k^*} = k - k^*$.

It must hold that $k - k^* \leq i_{k^*} - 1$ since the left hand side is the number of agents that took an item with strictly positive value before i_{k^*} arrived, and this number is at most $i_{k^*} - 1$, or else i_{k^*} would not be able to take an item with strictly positive value. We conclude that $i_k - i_{k^*} \leq i_{k^*} - 1$ or equivalently $i_k \leq 2i_{k^*} - 1$. Hence, the total value that agents $i_k, \ldots, i_{k^*}$ contribute is no more than

$$H_{i_k} - H_{i_{k^*}-1} \leq H_{2i_{k^*}-1} - H_{i_{k^*}-1} = \sum_{j=i_{k^*}}^{2i_{k^*}-1} \frac{1}{j} \leq i_{k^*} \cdot \frac{1}{i_{k^*}} = 1.$$

□

We now describe a prioritization algorithm using r priority classes that has a matching bound.

Theorem 7. *For arbitrary valuations, there exists an algorithm using $r \geq 1$ priority classes that achieves an approximation ratio of $\Omega(\frac{r}{\log n})$.*

Proof. Given a weighted graph G, fix some maximum matching M and discard any agents not in the matching. For any remaining agent $\ell \in L$, let $M(\ell)$ be the item matched to agent ℓ in the matching M. Let w_ℓ be $w(\ell, M(\ell))$. We define sets B_i where ℓ belongs to B_i if $2^i \leq w_\ell < 2^{i+1}$. Let $Z = \{B_{i_1}, B_{i_2}, \ldots, B_{i_r}\}$ be the collection of the r B_i's with the highest contribution to M, where $i_j > i_{j+1}$ for $j = 1, \ldots, r-1$.

The jth priority class, C_j is a random subset of B_{i_j} where every agent is taken with probability p. As the adversary is oblivious, it can be considered as though it predetermines a global order of agent arrivals and in particular the relative order of arrival for each of the sets B_{i_j}. In fact, only agents from C_j will arrive and the agents in $B_{i_j} \setminus C_j$ are plebeians and arrive last.

Re-index agents in $\cup_j B_{i_j}$ such that the agents in B_{i_1}, ordered by the adversary determined order of arrival, have indices $1, ..., n_1$. More generally, agents in B_{i_j}, ordered by the adversary determined order of arrival, have indices $n_{j-1}+1, ..., n_j$. For convenience, we define $n_0 = 0$, and define the maximum value of an empty set to be zero.

As done in Theorem 1, we say that item j is unavailable after event i, $0 \leq i \leq n_r$, if the item was chosen by one of the agents $1, ..., i$ from $\cup_j C_j$, before the arrival of agent $i+1$. For all $0 \leq i < k \leq n_r$ define n_k^i as follows:

$$n_k^i := \begin{cases} 1 & if\ M(k) \text{ is not available after event } i \\ 0 & else \end{cases}$$

Let $S_i^w := \sum\limits_{k=i+1}^{n_r} n_k^i \cdot w_k$. For $0 \leq i \leq n_r - 1$ the following holds

$$S_{i+1}^w \leq S_i^w - n_{i+1}^i \cdot w_{i+1} + I_{i+1} \cdot \max\{w_j | i+2 \leq j \leq n_r\}, \tag{5}$$

where $I_{i+1} = 1$ if agent $i+1$ is chosen to a priority class and otherwise $I_{i+1} = 0$.

The difference $S_{i+1}^w - S_i^w$ consists of several components. $S_{i+1}^w - S_i^w$ decreases by $n_{i+1}^i \cdot w_{i+1}$, and if $I_{i+1} = 1$ the difference may increase by $\max\{w_j | i+2 \leq j \leq n_r\}$ (this happens when agent $i+1$ takes some item in $\{M(i+1), \ldots, M(n_r)\}$).

By taking expectation over (5), using linearity of expectation, and noting that an agent takes an item (any item) with probability p, we get that

$$\mathbb{E}[S_{i+1}^w] \leq \mathbb{E}[S_i^w] - \mathbb{E}[n_{i+1}^i \cdot w_{i+1}] + p \cdot \max\{w_j | i+2 \leq j \leq n_r\},$$

or equivalently

$$\mathbb{E}[S_{i+1}^w] - \mathbb{E}[S_i^w] \leq -\mathbb{E}[n_{i+1}^i \cdot w_{i+1}] + p \cdot \max\{w_j | i+2 \leq j \leq n_r\}.$$

By taking the sum of i over 0 to $n_r - 1$ and noting that $\mathbb{E}[S_{n_r}^w] = \mathbb{E}[S_0^w] = 0$ we get that

$$\sum_{i=0}^{n_r-1} \mathbb{E}[n_{i+1}^i] \cdot w_{i+1} \leq p \cdot \sum_{i=0}^{n_r-1} \max\{w_j | i+2 \leq j \leq n_r\} = p \cdot \sum_{i=1}^{n_r} \max\{w_j | i+1 \leq j \leq n_r\}.$$

It now follows that

$$\sum_{i=0}^{n_r-1} \mathbb{E}[n_{i+1}^i] \cdot w_{i+1} \leq p \sum_{l=1}^{r} \sum_{i=n_{l-1}+1}^{n_l} \max\{w_j | i+1 \leq j \leq n_r\} \leq p \sum_{l=1}^{r} \sum_{i=n_{l-1}+1}^{n_l} 2w_i = 2 \sum_{i=1}^{n_r} w_i.$$

Thus, we get that

$$\sum_{i=0}^{n_r-1} \mathbb{E}[n_{i+1}^i] \cdot w_{i+1} \le 2p \sum_{i=1}^{n_r} w_i. \tag{6}$$

Let R_i be the weight of the matching after event i. The sequence R_i is [weakly] ascending. We now show that

$$R_{i+1} \ge R_i + I_{i+1} \cdot w_{i+1} \cdot \big(1 - n_{i+1}^i\big), \tag{7}$$

by the following case analysis:

- If $n_{i+1}^i = 1$ or $I_{i+1} = 0$ then (7) follows directly from monotonicity of R_i.
- Else, $I_{i+1} = 1$ and $n_{i+1}^i = 0$, i.e., $M(i+1)$ was available after event i and agent $i+1$ was admitted, hence the size of the matching increases by at least w_{i+1}, and indeed (7) is equivalent to $R_{i+1} \ge R_i + w_{i+1}$.

Now, taking the expectation over (7) we derive that

$$\mathbb{E}[R_{i+1}] \ge \mathbb{E}\Big[R_i + I_{i+1} \cdot w_{i+1} \cdot \big(1 - n_{i+1}^i\big)\Big].$$

It now follows from linearity of expectation and the fact that I_{i+1} is independent of n_{i+1}^i that

$$\mathbb{E}[R_{i+1}] \ge \mathbb{E}[R_i] + \mathbb{E}[I_{i+1}] w_{i+1} \Big(1 - \mathbb{E}[n_{i+1}^i]\Big),$$

or equivalently

$$\mathbb{E}[R_{i+1}] - \mathbb{E}[R_i] \ge p w_{i+1} \Big(1 - \mathbb{E}[n_{i+1}^i]\Big).$$

Taking the sum of i over $i = 0$ to $i = n_r - 1$ we get that

$$\mathbb{E}[R_{n_r}] - \mathbb{E}[R_0] \ge \sum_{i=0}^{n_r-1} p w_{i+1} - p \sum_{i=0}^{n_r-1} \mathbb{E}[n_{i+1}^i] w_{i+1}.$$

Note that $\mathbb{E}[R_0] = 0$. Using the bound for $\sum_{i=0}^{n_r-1} \mathbb{E}[n_{i+1}^i] w_{i+1}$ from (6) we have that

$$\mathbb{E}[R_{n_r}] \ge p \left(\sum_{i=0}^{n_r-1} w_{i+1} - 2p \sum_{i=0}^{n_r-1} w_{i+1} \right) = (p - 2p^2) \sum_{i=0}^{n_r-1} w_{i+1}.$$

To maximize the expected matching we take $p = \frac{1}{4}$ and we get

$$\mathbb{E}[R_{n_r}] \ge \frac{1}{8} \sum_{i=0}^{n_r-1} w_{i+1}.$$

Note that the total contribution of the top r classes to M is at least $\frac{r}{2 \log n}$ of the optimal social welfare, hence

$$\mathbb{E}[R_{n_r}] \ge \frac{r}{16 \log n} \mathrm{Opt},$$

where Opt is the weight of the maximum weighted matching. $\square$

5 Extention to q-capped Allocations

Up to now we have restricted agents to take at most one item. We now turn to agents with additive valuations and increase their quota to taking no more than q items each. We refer to such allocations as q-capped allocations. The proofs of the theorems in this section appear in Appendix A.

As in previous sections, we consider both inconsiderate strangers and friends.

5.1 Inconsiderate Strangers

Theorem 1 for unit-demand valuations extends to the case of additive valuations, with an additional loss of factor q.

Theorem 8. *For any $0 \leq \alpha \leq 1$, prioritizing every agent with probability $\frac{\alpha}{2q}$ gives an $\frac{\alpha}{4q}$ approximation to the size of the optimal q-capped allocation, for any unweighted graph with a maximum q-capped allocation of size $\geq \alpha n$.*

We also have a matching impossibility result.

Theorem 9. *For an unweighted graph with a maximum q-capped allocation of value αn, no prioritization algorithm can achieve an approximation to the maximum q-capped allocation greater than $O(\frac{\alpha}{q})$.*

5.2 Inconsiderate Friends

Theorem 7 for unit-demand valuations extends to the case of additive valuations, with an additional loss of factor q.

Theorem 10. *For arbitrary additive valuations, there exists an algorithm using $r \geq 1$ priority classes that achieves an approximation ratio of $\Omega(\frac{r}{q \log n})$ to the maximum q-capped allocation.*

We also have an impossibility result.

Theorem 11. *For arbitrary additive valuations, adding a business class (in addition to the default class) does not give an approximation ratio greater than $\max\{O(\frac{1}{q \log n}), \frac{1}{\sqrt{n}}\}$. Moreover, no algorithm that uses r priority classes can achieve an approximation ratio greater than $\max\{O(\frac{r}{q \log n}), \frac{1}{\sqrt{n}}\}$.*

6 Discussion

In this paper we study nearly-efficient allocation of goods to inconsiderate agents that arrive over time. Previous work on online resource allocation concentrate on either non-strategic agents or on money as a tool for creating appropriate incentives for driving the agents into desired outcomes. We consider settings in which agents are strategic and inconsiderate, yet money cannot be used to alleviate the problem. We propose a new class of algorithms, called prioritization algorithms, where agents are assigned to a small set of priority classes, and higher classes always precede lower ones (order within a given class is arbitrary). We show that simple prioritization algorithms can lead to approximately optimal welfare in various allocation settings, even when the entire inventory is free and agents behave selfishly and inconsiderately. We hope that prioritization algorithms can serve as a useful tool in additional online problems.

References

[1] Gagan Aggarwal, Gagan Goel, Chinmay Karande, and Aranyak Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 1253–1264, 2011.

[2] Itai Ashlagi, Felix A. Fischer, Ian A. Kash, and Ariel D. Procaccia. Mix and match. In *Proceedings 11th ACM Conference on Electronic Commerce (EC-2010), Cambridge, Massachusetts, USA, June 7-11, 2010*, pages 305–314, 2010.

[3] Haris Aziz, Hau Chan, Barton E. Lee, and David C. Parkes. Mechanism design without money for conmmon goods. 2018. arXiv:1806.00960v1 [cs.GT].

[4] Allan Borodin, Morten N. Nielsen, and Charles Rackoff. (incremental) priority algorithms. *Algorithmica*, 37(4):295–326, 2003.

[5] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[6] Vincent Cohen-Addad, Alon Eden, Michal Feldman, and Amos Fiat. The invisible hand of dynamic market pricing. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 383–400, 2016.

[7] Paul Duetting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 540–551, 2017.

[8] Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and S. Muthukrishnan. Online stochastic matching: Beating 1-1/e. In *Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '09, pages 117–126, Washington, DC, USA, 2009. IEEE Computer Society.

[9] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial walrasian equilibrium. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 61–70, 2013.

[10] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 123–135, 2015.

[11] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95 – 124, 1999.

[12] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and David C. Parkes. Adaptive limited-supply online auctions. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, EC '04, pages 71–80, New York, NY, USA, 2004. ACM.

[13] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. *STOC*, pages 352–358, 1990.

[14] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proceedings of the 2Nd ACM Conference on Electronic Commerce*, EC '00, pages 233–241, New York, NY, USA, 2000. ACM.

[15] Nina Mazar, Kristina Shampanier, and Dan Ariely. When retailing and las vegas meet: Probabilistic free price promotions. *Management Science*, 63(1):250–266, 2017.

[16] Aranyak Mehta. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4):265–368, 2013.

[17] David C. Parkes. Online mechanisms. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 16, pages 411–439. Cambridge University Press, New York, NY, USA, 2007.

[18] Ariel D. Procaccia. Cake cutting: not just child's play. *Commun. ACM*, 56(7):78–87, 2013.

[19] Ariel D. Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. *ACM Trans. Economics and Comput.*, 1(4):18:1–18:26, 2013.

[20] Kristina Shampanier, Nina Mazar, and Dan Ariely. Zero as a special price: The true value of free products. *Marketing Science*, 2007.

[21] David Robert Martin Thompson and Kevin Leyton-Brown. Dominant-strategy auction design for agents with uncertain, private values. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2011, San Francisco, California, USA, August 7-11, 2011*, 2011.

A Missing Proofs

A.1 Proof of Theorem 8

Given an unweighted graph G, fix an optimal q-capped allocation M. Index the agents by the adversary global order starting with agent 1. Let $M(k)$, $1 \leq k \leq n$, be the set of items allocated to agent k in M (recall $|M(k)| \leq q$ for each $k \in L$). Denote $M(k) := \{m_k^1, ..., m_k^{q_k}\}$. Note that we do not know the adversary global order, nor do we know the $M(k)$'s.

For the purpose of analysis imagine that all agents arrive in the adversary global order where the ith event is the arrival of agent i. Agents are either allowed to make a choice or assigned to the Plebeian class (and thus delayed). This assignment to the Plebeian class is done on the fly.

We say that item j is unavailable after event i, $0 \leq i \leq n$ if it was taken by a prioritized agent j s.t. $j \in \{1, \ldots, i\}$.

For $0 \leq i < k \leq n$ and $1 \leq j \leq q_k$ define $n_{k,j}^i$ as follows:

$$n_{k,j}^i := \begin{cases} 1 & if\ m_k^j \text{ is unavailable after event } i \\ 0 & else. \end{cases}$$

Let $N_k^i := \sum_{j=1}^{m_k^{q_k}} n_{k,j}^i$ and $S_i := \sum_{k>i} N_k^i$. The following holds:

$$S_{i+1} \leq S_i - N_{i+1}^i + I_{i+1} \cdot q, \tag{8}$$

where $I_{i+1} = 1$ if agent $i+1$ is prioritized and otherwise $I_{i+1} = 0$.

The difference $S_{i+1} - S_i$ consists of several components. $S_{i+1} - S_i$ clearly decreases by N_{i+1}^i, and if $I_{i+1} = 1$ this difference may increase by q — this happens when agent $i+1$ takes q items from $\bigcup_{j=i+2}^{n} M(j)$.

Taking the expectation over (8), using linearity of expectation, and noting that agent $i+1$ is prioritized with probability $\frac{\alpha}{2q}$, we get that

$$\mathbb{E}[S_{i+1}] - \mathbb{E}[S_i] \leq -\mathbb{E}[N_{i+1}^i] + \frac{\alpha}{2}.$$

Taking the sum of i from 0 to $n-1$, we get that:

$$\mathbb{E}[S_n] - \mathbb{E}[S_0] \leq -\sum_{i=0}^{n-1} \mathbb{E}[N_{i+1}^i] + n \cdot \frac{\alpha}{2}.$$

Note that $\mathbb{E}[S_n] = \mathbb{E}[S_0] = 0$ and hence:

$$\sum_{i=0}^{n-1} \mathbb{E}[N_{i+1}^i] \leq n \cdot \frac{\alpha}{2}. \tag{9}$$

Let R_i be the value of the allocation (i.e., number of allocated items of value 1) after event i. The sequence R_i is [weakly] ascending. We now show that

$$R_{i+1} \geq R_i + I_{i+1} \cdot \big(|M(i+1)| - N_{i+1}^i\big), \tag{10}$$

by the following case analysis

- If $I_{i+1} = 0$ or $N^i_{i+1} = |M(i+1)|$ then (10) follows directly from monotonicity of R_i.
- The only remaining case is when $I_{i+1} = 0$ and $N^i_{i+1} < |M(i+1)|$, and then the size of the matching increases by at least the number of available valuable items from $M(i+1)$.

Taking the expectation over (10) and using linearity of expectation combined with the fact that I_{i+1} is independent of N^i_{i+1} we get that

$$\mathbb{E}[R_{i+1}] \geq \mathbb{E}[R_i] + \mathbb{E}[I_{i+1}]\Big(|M(i+1)| - \mathbb{E}[N^i_{i+1}]\Big),$$

or equivalently

$$\mathbb{E}[R_{i+1}] - \mathbb{E}[R_i] \geq \frac{\alpha}{2q}\Big(|M(i+1)| - \mathbb{E}[N^i_{i+1}]\Big).$$

Taking the sum for i from 0 to $n-1$, we get that

$$\mathbb{E}[R_n] = \mathbb{E}[R_n] - \mathbb{E}[R_0] \geq \frac{\alpha}{2q} \cdot n\alpha - \frac{\alpha}{2q}\sum_{i=0}^{n-1}\mathbb{E}[N^i_{i+1}] \geq \frac{n\alpha^2}{2q} - \frac{\alpha}{2q} \cdot \frac{n\alpha}{2} = \frac{n\alpha^2}{4q}.$$

Where the first equality follows from $\mathbb{E}[R_0] = 0$ and the last inequality follows from (9). This implies a bound on the approximation ratio of the algorithm:

$$\frac{\text{Alg}}{\text{Opt}} \geq \frac{n\alpha^2/4q}{n\alpha} = \frac{\alpha}{4q}.$$

where Opt is the size of the maximum weighted matching and Alg is the expected size of the matching achieved by the algorithm.

A.2 Proof of Theorem 9

Consider an instance with n agents and n items as depicted in Figure 2. As agents break ties (amongst zero valued items) from top to bottom, the first $\lceil \frac{\alpha n}{q} \rceil$ agents to arrive choose items $r_1, \ldots, r_{\alpha n}$ so there is no point in allowing more than $\lceil \frac{\alpha n}{q} \rceil$ agents to take items. The adversary can schedule αn random slots for agents $\ell_1, \ldots, \ell_{\alpha n}$. The best way to choose a priority class, in this case, is choosing a subset of $\lceil \frac{\alpha n}{q} \rceil$ agents randomly, resulting in an expected social welfare of size $\theta\big(\frac{\alpha^2 n}{q}\big)$. Hence, the approximation ratio is $O\big(\frac{\alpha}{q}\big)$.

A.3 Proof of Theorem 10

Given a weighted graph G, fix some optimal q-capped allocation M and exclude any agent that does not contribute a positive value. For any remaining agent $\ell \in L$, let $M(\ell) := \{m^1_\ell, \ldots, m^{q_i}_\ell\}$ be the subset allocated to agent ℓ in M. Let $w^1_\ell, \ldots, w^{q_\ell}_\ell$ be the corresponding weights of $M(\ell)$. Denote $w_\ell = \sum_{j=1}^{q_\ell} w^j_\ell$. We define sets B_i where ℓ belongs to B_i if $2^i \leq w_\ell < 2^{i+1}$. Let $Z = \{B_{i_1}, B_{i_2}, \ldots, B_{i_r}\}$ be the collection of the r B_i's with the highest contribution to M, where $i_j > i_{j+1}$ for $j = 1, \ldots, r-1$.

The jth priority class, C_j is a random subset of B_{i_j} where every agent is taken with probability p, that will be defined later. As the adversary is oblivious, it can be considered as though it predetermines a global order of agent arrivals and in particular the relative order of arrival for each of the sets B_{i_j}. In fact, only agents from C_j will arrive and the agents in $B_{i_j} \setminus C_j$ are plebeians and arrive last.

Re-index agents in $\cup_j B_{i_j}$ such that the agents in B_{i_1}, ordered by the adversary determined order of arrival, have indices $1, ..., n_1$. More generally, agents in B_{i_j}, ordered by the adversary determined order of arrival, have indices $n_{j-1}+1, ..., n_j$. For convenience, we define $n_0 = 0$.

As done in Theorem 1, we say that item j is unavailable after event i, $0 \le i \le n_r$, if the item was chosen by one of the agents $1, ..., i$ from $\cup_j C_j$, before the arrival of agent $i+1$. For each $0 \le i < k \le n_r$ and $1 \le j \le q_k$ define $n^i_{k,j}$ as follows:

$$n^i_{k,j} := \begin{cases} 1 & if\ m^j_k \text{ is not available after event } i \\ 0 & else \end{cases}$$

Let $N^i_k := \sum_{j=1}^{m^{q_k}_k} n^i_{k,j} \cdot w^j_k$, and let $S^w_i := \sum_{k=i+1}^{n_r} N^i_k$. For a set of real numbers A let $\max(q)A$ be the sum of the $min\{q, |A|\}$ highest values in A. If A is empty we define $\max(q)A = 0$. For $0 \le i \le n_r - 1$ the following holds:

$$S^w_{i+1} \le S^w_i - N^i_{i+1} + I_{i+1} \cdot \max(q)\{w^l_j | i+2 \le j \le n_r, 1 \le l \le q_j\} \tag{11}$$

where $I_{i+1} = 1$ if agent $i+1$ is chosen to a priority class and otherwise $I_{i+1} = 0$. The difference $S^w_{i+1} - S^w_i$ consists of several components. $S^w_{i+1} - S^w_i$ is clearly decreases by N^i_{i+1}, and if $I_{i+1} = 1$ the difference may increase by $\max(q)\{w^l_j | i+2 \le j \le n_r, 1 \le l \le q_j\}$, this happens when agent $i+1$ takes q items from $\bigcup_{i+2 \le k \le n_r} M(k)$. Taking the expectation over (11), using linearity of expectation and noting that an agent is allowed to take items with probability p, we get that

$$\mathbb{E}[S^w_{i+1}] - \mathbb{E}[S^w_i] \le \mathbb{E}[N^i_{i+1}] + p \cdot \max(q)\{w^l_j | i+2 \le j \le n_r, 1 \le l \le q_j\}.$$

Taking the sum of i from 0 to $n_r - 1$ and noting that $\mathbb{E}[S^w_{n_r}] = \mathbb{E}[S^w_0] = 0$ we get

$$\begin{aligned}
\sum_{i=0}^{n_r-1} \mathbb{E}[N^i_{i+1}] &\le p \cdot \sum_{i=0}^{n_r-1} \max(q)\{w^l_j | i+2 \le j \le n_r, 1 \le l \le q_j\} \\
&= p \sum_{l=1}^{r} \sum_{i=n_{l-1}+1}^{n_l} \max(q)\{w^l_j | i+2 \le j \le n_r, 1 \le l \le q_j\} \\
&\le p \sum_{l=1}^{r} \sum_{i=n_{l-1}+1}^{n_l} q \cdot 2w_i = 2pq \sum_{i=1}^{n_r} w_i.
\end{aligned}$$

It follows that

$$\sum_{i=0}^{n_r-1} \mathbb{E}[N^i_{i+1}] \le 2pq \sum_{i=1}^{n_r} w_i. \tag{12}$$

Define R_i as the total weight of the partial allocation after event i. The sequence R_i is [weakly] ascending. We show that

$$R_{i+1} \geq R_i + I_{i+1} \left(\sum_{j=1}^{m_{i+1}^{q_{i+1}}} (1 - n_{i+1,j}^i) w_{i+1}^j \right), \tag{13}$$

by the following case analysis:

- if $I_{1+1} = 0$ then (13) follows directly from monotonicity of R_i.
- else, the the contribution of agent $i+1$ to the allocation is at least the total value of the available items from $M(i+1)$.

Now, taking the expectation over (13), and using linearity of expectation combined with the fact that I_{i+1} is independent of $n_{i+1,j}^i$ we get

$$\mathbb{E}[R_{i+1}] - \mathbb{E}[R_i] \geq \mathbb{E}[I_{i+1}] \Big(\sum_{j=1}^{m_{i+1}^{q_{i+1}}} (1 - \mathbb{E}[n_{i+1,j}^i]) w_{i+1}^j \Big),$$

or equivalently

$$\mathbb{E}[R_{i+1}] - \mathbb{E}[R_i] \geq p w_{i+1} - p \mathbb{E}[N_{i+1,j}^i].$$

Taking the over i from $i = 0$ to $i = n_r - 1$ we get

$$\mathbb{E}[R_{n_r}] - \mathbb{E}[R_0] \geq p \sum_{i=0}^{n_r - 1} w_{i+1} - p \sum_{i=0}^{n_r - 1} \mathbb{E}[N_{i+1,j}^i].$$

Note that $\mathbb{E}[R_0] = 0$. Using the bound from (12) we have that

$$\mathbb{E}[R_{n_r}] \geq p \Big(\sum_{i=0}^{n_r-1} w_{i+1} - 2p^2 q \sum_{i=0}^{n_r-1} w_{i+1} \Big) = (p - 2p^2 q) \sum_{i=0}^{n_r-1} w_{i+1}.$$

To maximize the expected value of the allocation we take $p = \frac{1}{4q}$ and get

$$\mathbb{E}[R_{n_r}] \geq \frac{1}{8} \sum_{i=0}^{n_r-1} w_{i+1}$$

Note that $\sum_{i=0}^{n_r-1} w_{i+1} \geq \frac{r}{\log n} \frac{\mathrm{Opt}}{2}$ so

$$\mathbb{E}[R_{n_r}] \geq \frac{r}{16 \log n} \mathrm{Opt}$$

where Opt is the weight of the maximum q-capped allocation.

A.4 Proof of Theorem 11

First consider the case where $q \geq \sqrt{n}$. Let $G = (L \cup R, E)$ be a graph depicted in Figure 5 with n agents, $\ell_1, \ldots, \ell_n$, and n items, $r_1, \ldots, r_n$. The set of edges is $E = \{(\ell_i, r_i) |\, 1 \leq i \leq n\}$, and for all $1 \leq i \leq n$ $w(\ell_i, r_i) = 1$. Clearly the optimal q-capped allocation is of value n (agent ℓ_i takes r_i). After $\left\lceil \frac{n}{\sqrt{n}} \right\rceil = \lceil \sqrt{n} \rceil$ agents arrive there is no item left. Allowing any subset of $\lceil \sqrt{n} \rceil$ agents to take items results in an allocation of value no greater than $\lceil \sqrt{n} \rceil$, hence an approximation ratio of size $O(1/\sqrt{n})$.

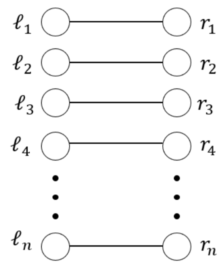

Figure 5: Agents on the left and items on the right.

Next, consider $q < \sqrt{n}$. For $q = 1$, see Theorem 6. For $q \geq 2$, consider the bipartite graph depicted in Figure 6, with n agents, $\ell_1, \ldots, \ell_n$, and n items, $r_1, \ldots, r_n$. The set of edges is $E = \{(\ell_i, r_j) |\, q \leq i \leq n\}$, and for all $q \leq i \leq n$ $w(\ell_i, r_i) = \frac{1}{i}$.

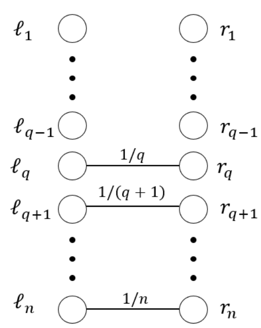

Figure 6: Agents on the left and items on the right. The number that appears on the edge represents the weight of that edge.

Clearly, in the maximum matching agent ℓ_i, $q \leq i$, is matched to item r_i, resulting in a total weight of

$$H_n - H_{q-1} = \sum_{i=q}^{n} \frac{1}{i} \approx \ln n - \ln(q-1) \approx \ln n.$$

Fix an arbitrary subset of agents. We describe an adversary global order and tie breaking rule such that the resulting social welfare of this subset is at most $O\big(\frac{1}{q}\big)$. This implies the theorem, since for a single priority class (i.e., partitioning the agents to two subsets) we get at most $O\big(\frac{1}{q \ln n}\big)$ of the optimal welfare, and with r priority classes we get at most $O\big(\frac{r}{q \ln n}\big)$ of the optimal welfare.

The global ordering of the agents and tie breaking rule is as follows: agents are ordered from high indices to low indices (bottom to top in figure), and always break ties among zero valued items in the favor of items with lower indices (higher items in figure). Consider a subset of agents that forms a priority class $C = \{i_1, ..., i_k\}$, where $i_j < i_{j+1}$ for $j = 1, \ldots, k-1$. Clearly, there is some index, i_{k^*}, such that for all $j \geq k^*$ agent i_j takes one item of positive value among the bundle she takes; and for all $j < k^*$ agent i_j takes only zero valued items. Assume without loss of generality that $i_{k^*}, \ldots, i_k$ are consecutive agents (i.e., when for all $k^* \leq j \leq k-1$, $i_j + 1 = i_{j+1}$). This is without loss since agents with lower indices contribute higher values. Therefore, $i_k - i_{k^*} = k - k^*$. It must hold that $(q-1)(k-k^*) \leq i_{k^*} - 1$ since the left hand side is the number of items that were taken by agents before i_{k^*} arrived, and it is at most $i_{k^*} - 1$, otherwise i_{k^*} would not be able to take an item with strictly positive value. We conclude $(q-1)(i_k - i_{k^*}) \leq i_{k^*} - 1$ or equivalently $i_k \leq \frac{qi_{k^*}-1}{q-1}$. Hence, the total value of agents $i_k, \ldots, i_{k^*}$ is not greater than

$$\begin{aligned}
H_{i_k} - H_{i_{k^*}-1} &\leq H_{\left\lfloor \frac{qi_{k^*}-1}{q-1} \right\rfloor} - H_{i_{k^*}-1} \\
&= \sum_{j=i_{k^*}}^{\left\lfloor \frac{qi_{k^*}-1}{q-1} \right\rfloor} \frac{1}{j} \\
&\leq \left(\left\lfloor \frac{qi_{k^*}-1}{q-1} \right\rfloor - i_{k^*} + 1 \right) \cdot \frac{1}{i_{k^*}} \\
&\leq \left(\frac{qi_{k^*}-1}{q-1} - i_{k^*} + 1 \right) \cdot \frac{1}{i_{k^*}} \\
&= \left(\frac{i_{k^*}}{q-1} - \frac{1}{q-1} + 1 \right) \cdot \frac{1}{i_{k^*}} \\
&= \frac{1}{q-1} + \frac{q-2}{i_{k^*}(q-1)}
\end{aligned}$$

if $i_{k^*} \geq q$ then the expression above is not greater than

$$\frac{1}{q-1} + \frac{q-2}{q(q-1)} \leq \frac{2}{q-1},$$

and if $i_{k^*} < q$, i_{k^*} would be able to take a positive value item only when $k = k^*$, hence, only a single agent takes a positive value item. The maximum value of an edge in the graph is $\frac{1}{q}$, which bounds the total contribution of the class.