

Social Welfare in One-sided Matchings: Random Priority and Beyond^{*}

Aris Filos-Ratsikas¹, Søren Kristoffer Stiil Frederiksen¹, and Jie Zhang²

¹ Department of Computer Science, Aarhus University
`{filosra,ssf}@cs.au.dk`

² Department of Computer Science, University of Oxford
`jie.zhang@cs.ox.ac.uk`

Abstract. We study the problem of approximate social welfare maximization (without money) in one-sided matching problems when agents have unrestricted cardinal preferences over a finite set of items. *Random priority* is a very well-known *truthful-in-expectation* mechanism for the problem. We prove that the approximation ratio of random priority is $\Theta(n^{-1/2})$ while no truthful-in-expectation mechanism can achieve an approximation ratio better than $O(n^{-1/2})$, where n is the number of agents and items. Furthermore, we prove that the approximation ratio of all ordinal (not necessarily truthful-in-expectation) mechanisms is upper bounded by $O(n^{-1/2})$, indicating that random priority is asymptotically the best truthful-in-expectation mechanism and the best ordinal mechanism for the problem.

1 Introduction

We study the problem of approximate social welfare maximization (without money) in one-sided matching problems when agents have unrestricted cardinal preferences over a finite set of items. Specifically, each agent has a valuation function mapping items to real numbers, which can be arbitrary. These valuation functions should be interpreted as von Neumann-Morgenstern utility functions, i.e. they induce orderings on outcomes, which are standardly defined up to positive affine transformations (multiplication by a positive scalar and shift by a scalar).

A (direct revelation) mechanism (without money) is a function J mapping vectors of valuation functions (valuation profiles) to matchings, that is, allocations of items to agents such that each agent is allocated exactly one item. Mechanisms can also be randomized, and then the function J is a random map.

^{*} The authors acknowledge support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, within which this work was performed. The authors also acknowledge support from the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council. Jie Zhang also acknowledges support from ERC Advanced Grant 321171 (ALGAME).

We will be interested in *truthful* mechanisms, that is, mechanisms that do not give incentives to agents to misreport their valuation functions. For randomized mechanisms, we are interested in mechanisms that are *truthful-in-expectation*, which means that no agent can increase its expected utility by misreporting.

A very intuitive and well-studied truthful-in-expectation mechanism is *random priority* (often also referred to as *random serial dictatorship*), which first fixes a uniformly random ordering of the agents and then serially matches each agent to its most preferred item from the set of still unmatched items, based on that ordering. Random priority is an *ordinal* mechanism, i.e., a mechanism that only depends on the ordering of items induced by the valuation functions and not the actual numerical values.

We will look to maximize the *social welfare*, that is, the sum of agents' valuations for the items they are matched with in the outcome of the mechanism. We measure the performance of a mechanism by its approximation ratio, which is the worst ratio between the (expected) social welfare of the mechanism and the welfare of the optimal allocation, over all valuation profiles. It is easy to see that the mechanism that always outputs the optimal outcome is not truthful.

For a meaningful discussion on social welfare maximization for von Neumann-Morgenstern utilities, one has to fix a canonical representation of the valuation functions [7, 9, 11, 12]. These functions are usually represented in one of two canonical forms, *unit-sum* [7, 12] (the valuations sum up to one) or *unit-range* [4, 11, 19] (all valuations are in the interval $[0, 1]$ with both 0 and 1 in the image of the function). Our main result is given by the following theorem, which holds for both normalizations.

Theorem 1. *The approximation ratio of random priority is $\Theta(1/\sqrt{n})$. Furthermore, random priority is asymptotically the best truthful-in-expectation mechanism and the best ordinal (not necessarily truthful-in-expectation) mechanism for the problem.*

The theorem also holds for an extension to the unit-range representation, when 0 is not required to be in the image of the function; we discuss how in Section 5. An implication of Theorem 1 is that the well-known *Probabilistic Serial* mechanism [6], which is often preferred to random priority, does not give better social welfare guarantees, even if we assume truthful reporting.

1.1 Discussion and related work

In the presence of incentives, the one-sided matching problem (often referred to as *the assignment problem* or *house allocation problem*) was originally defined in [13] and has been studied extensively ever since [6, 9, 12, 14, 18, 19]. There are several surveys discussing the problem [1, 17] and we refer the interested reader to those for a more detailed exposition. Random priority is a folklore mechanism that solves the problem fairly (in the sense of anonymity) and satisfies some additional nice properties; it is truthful-in-expectation and *ex-post Pareto efficient*. On the other hand, it is not *ex-ante Pareto efficient*. Most of the previous work

in literature [6, 13, 19] has mainly been directed towards designing mechanisms with desired properties that achieve efficiency criteria for different Pareto efficiency notions and less towards whether truthful mechanisms achieve “good levels” of social welfare. This is a really important question to ask, especially since the notion of approximation ratio gives us a systematic way of comparing mechanisms or proving their limitations. The term *approximate mechanism design without money* was used in [16] to describe problems where the goal is to approximately optimize some objective function, given the constraint of truthfulness. This approach has been adopted by a large body of computer science literature [3, 5, 9, 11, 12, 15] and the approximation ratio is now considered to be the predominant measure of efficiency for truthful mechanisms.

Social welfare maximization is arguably a less natural objective when agents are endowed with von Neumann-Morgenstern utilities, because of adding up valuations after normalization. On the other hand, it is quite widespread in *quasi-linear* settings. We strongly believe that considering the social welfare objective for von Neumann-Morgenstern utilities is just as natural and in fact there is a growing amount of literature that embraces the same idea and provides arguments to support it [5, 7, 11, 12].

A different approach, often encountered in literature, is to consider ordinal measures of efficiency. For example, Bhargat et al [5] calculate the approximation ratio of random priority when the objective is the maximization of *ordinal* social welfare, a notion of welfare that they define based solely on ordinal information. Ordinal measures of efficiency have also been studied in terms of incentives and approximation ratios [8, 15]. However, these measures do not encapsulate the “socially desired” outcome in the way that social welfare does, i.e., they do not necessarily maximize the aggregate happiness of individuals [3]. This is even more evident if one considers that the standard assumption in social choice and economics theory is that such an underlying cardinal structure exists, even if agents are not asked to report it. In our setting, not reporting the full cardinal information corresponds to using ordinal mechanisms, which by our main theorem, is enough for achieving the (asymptotically) best approximation guarantees. On the other hand, cardinal reports are also often encountered in literature, with the pseudo-market mechanism of [13] being a prominent example. Several cardinal mechanisms were also presented in [11, 12] for social welfare maximization and [10] for information elicitation. Our main theorem does not only prove the capabilities of random-priority but also the limitations of all truthful (including cardinal) mechanisms. In the full version of our paper, we give an example of a non-ordinal mechanism that actually provides better approximation guarantees for concrete sizes of the set of agents.

Independently of our work, Adamczyk et al. [2] study truthful mechanisms for social welfare in one-sided matchings when agents have von Neumann - Morgenstern utilities, normalized in the unit interval, but not necessarily unit-sum or unit-range. Their main result on the approximation ratio of random priority can be combined with some additional arguments to obtain our lower bounds but our upper bounds are more general.

2 Preliminaries

Let $N = \{1, \dots, n\}$ be a finite set of agents and $M = \{1, \dots, n\}$ be a finite set of indivisible items. An *outcome* is a matching of agents to items, that is, an assignment of items to agents where each agent gets assigned exactly one item. We can view an outcome μ as a vector $(\mu_1, \mu_2, \dots, \mu_n)$ where μ_i is the unique item matched with agent i . Let O be the set of all outcomes. Each agent i has a private valuation function mapping outcomes to real numbers that can be arbitrary except for two conditions; agents are indifferent between outcomes that match them to the same item and they are not indifferent between outcomes that match them to different items. The first condition implies that agents only need to specify their valuations for items instead of outcomes and hence the valuation function of an agent i can be instead defined as a map $u_i : M \rightarrow \mathbb{R}$ from items to real numbers. The second condition requires that valuation functions are injective, i.e., they induce a total ordering on the items. This is mainly for convenience, to avoid having to specify tie-breaking rules. As we discuss in Section 5, all of our results extend to most natural tie-breaking rules. Valuation functions are standardly considered to be well-defined up to positive affine transformations, that is, for item $j : j \rightarrow \alpha u_i(j) + \beta$ is considered to be a different representation of u_i . The two standard ways to fix the canonical representation of u_i in literature are *unit-range*, i.e., $\max_j u_i(j) = 1$ and $\min_j u_i(j) = 0$ and *unit-sum*, that is $\sum_j u_i(j) = 1$.

Let V be the set of all canonically represented valuation functions of an agent. Call $\mathbf{u} = (u_1, u_2, \dots, u_n)$ a *valuation profile* and let V^n be the set of all valuation profiles with n agents. A *direct revelation mechanism* (without money) is a function $J : V^n \rightarrow O$ mapping *reported* valuation profiles to matchings. For a randomized mechanism, we define J to be random map $J : V^n \rightarrow O$. Let $J(\mathbf{u})_i$ denote the restriction of the outcome of the mechanism to the i 'th coordinate, which is the item assigned to agent i by the mechanism.

We will be interested in *truthful mechanisms*, that is, mechanisms that do not incentivize agents to report anything other than their true valuation functions. Formally, a mechanism J is truthful if for each agent i and all $\mathbf{u} = (u_i, u_{-i}) \in V^n$ and $\tilde{u}_i \in V$ it holds that $u_i(J(u_i, u_{-i})_i) \geq u_i(J(\tilde{u}_i, u_{-i})_i)$, where u_{-i} denotes the valuation profile \mathbf{u} without the i 'th coordinate. In other words, if u_i is agent i 's true valuation function, then it has no incentive to misreport. For randomized mechanisms, a mechanism is *truthful-in-expectation* if for each agent i and all $\mathbf{u} = (u_i, u_{-i}) \in V^n$ and $\tilde{u}_i \in V$ it holds that $\mathbb{E}[u_i(J(u_i, u_{-i})_i)] \geq \mathbb{E}[u_i(J(\tilde{u}_i, u_{-i})_i)]$.

A class of mechanisms that turns out to be important for our purposes is that of neutral and anonymous mechanisms. Formally, a mechanism is *anonymous* if for any valuation profile (u_1, u_2, \dots, u_n) , every agent i and any permutation $\pi : N \rightarrow N$ it holds that $J(u_1, u_2, \dots, u_n)_i = J(u_{\pi(1)}, u_{\pi(2)}, \dots, u_{\pi(n)})_{\pi(i)}$. By this definition, in an anonymous mechanism, agents with exactly the same valuation functions must have the same probabilities of receiving each item. Similarly, a mechanism is *neutral* if for any valuation profile (u_1, u_2, \dots, u_n) , every item j and any permutation $\sigma : M \rightarrow M$ it holds that $J(u_1, u_2, \dots, u_n)_i = \sigma^{-1}(J(u_1 \circ \sigma, u_2 \circ \sigma, \dots, u_n \circ \sigma)_i)$, i.e., the mechanism is invariant to the indices of the items.

We will consider both ordinal and cardinal mechanisms. A mechanism J is *ordinal* if for any i , any valuation profile $\mathbf{u} = (u_i, u_{-i})$ and any valuation function u'_i such that for all $j, j' \in M$, $u_i(j) < u_i(j') \Leftrightarrow u'_i(j) < u'_i(j')$, it holds that $J(u_i, u_{-i}) = J(u'_i, u_{-i})$. A mechanism for which the above does not necessarily hold is *cardinal*. Informally, ordinal mechanisms operate solely based on the *ordering* of items induced by the valuation functions and not the actual numerical values themselves.

We measure the performance of a mechanism by its approximation ratio,

$$ar(J) = \inf_{\mathbf{u} \in V^n} \frac{\sum_{i=1}^n u_i(J(\mathbf{u})_i)}{\max_{\mu \in O} \sum_{i=1}^n u_i(\mu_i)}.$$

The quantity $\sum_{i=1}^n u_i(J(\mathbf{u})_i)$ is the *social welfare* of mechanism J on valuation profile \mathbf{u} and $\max_{\mu \in O} \sum_{i=1}^n u_i(\mu_i)$ is the social welfare of the optimal matching. For ease of notation, let $w^*(\mathbf{u}) = \max_{\mu \in O} \sum_{i=1}^n u_i(\mu_i)$. For randomized mechanisms, we are interested in the *expected social welfare* $\mathbb{E}[\sum_{i=1}^n u_i(J(\mathbf{u})_i)]$ of mechanism J and the approximation ratio is defined accordingly.

Next we will state a lemma that will be useful for our proofs. These kinds of lemmas are standard in literature (e.g. see [11, 12]). The lemma implies that when trying to prove upper bounds on the approximation ratio of mechanisms, it suffices to consider mechanisms that are anonymous.

Lemma 1. *For any mechanism J , there exists an anonymous mechanism J' such that $ar(J') \geq ar(J)$. Furthermore, if J is truthful (for deterministic mechanisms) or truthful-in-expectation (for randomized mechanisms) then it holds that J' is truthful-in-expectation.*

We will particularly be interested in the mechanism *random priority*. Random priority fixes an ordering of the agents uniformly at random and then lets them pick their favorite items from the set of available items based on this ordering. Note that random priority is truthful-in-expectation, ordinal, anonymous and neutral. We conclude the section with the following lemma.

Lemma 2. *For any valuation profile \mathbf{u} , the optimal allocation on \mathbf{u} is a possible outcome of random priority.*

3 Unit-range valuation functions

In this section, we assume that the representation of the valuation functions is unit-range. It will be useful to consider a special class of valuation functions C_ϵ that we will refer to as *quasi-combinatorial valuation functions*, a straightforward adaptation of a similar notion in [11]. Informally, a valuation function is quasi-combinatorial if the valuations of each agent for every item are close to 1 or close to 0 (the proximity depends on ϵ). Formally,

$$C_\epsilon = \{u \in V \mid u(M) \subset [0, \epsilon) \cup (1 - \epsilon, 1]\},$$

where $u(M)$ is the image of the valuation function u . Let $C_\epsilon^n \subseteq V^n$ be the set of all valuation profiles with n agents whose valuation functions are in C_ϵ . The following lemma implies that when we are trying to prove a lower bound on the approximation ratio of random priority, it suffices to restrict our attention to quasi-combinatorial valuation profiles $C_\epsilon^n \subseteq V^n$ for any value of ϵ .

Lemma 3. *Let J be an ordinal, anonymous and neutral randomized mechanism for unit-range representation, and let $\epsilon > 0$. Then*

$$ar(J) = \inf_{\mathbf{u} \in C_\epsilon^n} \frac{\mathbb{E}[\sum_{i=1}^n u_i(J(\mathbf{u})_i)]}{w^*(\mathbf{u})}.$$

The lemma formalizes the intuition that because the mechanism is ordinal, the worst-case approximation ratio is encountered on extreme valuation profiles. In fact, the proof of the lemma inductively “pushes” agents’ valuations towards 0 or towards 1, without increasing the approximation ratio each time and ends up with a quasi-combinatorial profile.

For unit-range representation, Theorem 1 is given by the following lemmas.

Lemma 4. *For unit-range representation, $ar(RP) = \Omega(n^{-1/2})$.*

Proof. Because of Lemma 3, for computing a lower bound on the approximation ratio of random priority, it is sufficient to only consider quasi-combinatorial valuation profiles. Let $\epsilon \leq 1/n^3$. Then, there exists $k \in \mathbb{N}$ such that

$$|k - w^*(\mathbf{u})| \leq \frac{1}{n^2},$$

where $w^*(\mathbf{u})$ is the social welfare of the maximum weight matching on valuation profile \mathbf{u} . Since random priority can trivially achieve an expected welfare of 1 (for any permutation the first agent will be matched to its favorite item), we can assume that $k \geq \sqrt{n}$, otherwise we are done. Note that the maximum weight matching $\mu^* \in O$ assigns k items to agents with $u_i(\mu_i) \in (1 - \epsilon, 1]$. Since random priority is anonymous and neutral, without loss of generality we can assume that these agents are $\{1, \dots, k\}$ and for every agent $j \in N$, it holds that $\mu_j^* = j$. Thus $u_j(j) \in (1 - \epsilon, 1]$ for $j = 1, \dots, k$ and $u_j(j) \in [0, \epsilon)$ for $j = k + 1, \dots, n$.

Consider any run of random priority; one agent is selected in each round. Let $l \in \{0, \dots, n - 1\}$ be any of the n rounds. We will now define the following sets:

$$U_l = \{j \in \{1, \dots, n\} \mid \text{agent } j \text{ has not been selected prior to round } l\}$$

$$G_l = \{j \in U_l \mid u_j(j) \in (1 - \epsilon, 1] \text{ and item } j \text{ is still unmatched}\}$$

$$B_l = \{j \in U_l \mid u_j(j) \in [0, \epsilon) \text{ or item } j \text{ has already been matched to some agent}\}$$

These three families of sets should be interpreted as three sets that change over the course of the execution of random priority. U_l is the set of agents yet to be matched, which is partitioned into G_l , the set of “good” agents, that guarantee a welfare of almost 1 when picked, and B_l , the set of “bad” agents, that do not guarantee any contribution to the social welfare. For the purpose of calculating

a lower bound, we will simply bound the sizes of the sets in these families. Obviously, $G_0 = \{1, \dots, k\}$ and $B_0 = \{k+1, \dots, n\}$.

The probability that an agent $i \in G_l$ is picked in round l of random priority is $|G_l|/(|G_l| + |B_l|)$, whereas the probability that an agent $i \in B_l$ is picked is $|B_l|/(|G_l| + |B_l|)$. By the discussion above, we can assume that whenever an agent from G_l is picked its contribution to the social welfare is at least $1 - \epsilon$ whereas the contribution from an agent picked from B_l is less than ϵ . In other words, the expected contribution to the social welfare from round l is at least $|G_l|/(|G_l| + |B_l|) - \epsilon$.

We will now upper bound $|G_l|$ and lower bound $|B_l|$ for each l . Consider round l and sizes $|G_l|$ and $|B_l|$. First suppose that some agent i from G_l is picked and the agent is matched with item j . If $j \neq i$ and agent j is in G_l , then $|G_{l+1}| = |G_l| - 2$ and $|B_{l+1}| = |B_l| + 1$, since agent j no longer has its item from the optimal allocation available and so agent j is in B_{l+1} . On the other hand, if $j = i$ or agent j is in B_l then $|G_{l+1}| = |G_l| - 1$ and $|B_{l+1}| = |B_l|$. In either case, $|G_{l+1}| \geq |G_l| - 2$ and $|B_{l+1}| \leq |B_l| + 1$. Intuitively, the picked agent might take away some item from a good agent and turn it into a bad agent.

Now suppose that agent i from B_l is picked and the agent is matched with item j . If agent j is in G_l then $|G_{l+1}| = |G_l| - 1$ and $|B_{l+1}| = |B_l|$, since agent j no longer has its item from the optimal allocation available and so agent j is in B_{l+1} . On the other hand, if agent j is in B_l then $|G_{l+1}| = |G_l|$ and $|B_{l+1}| = |B_l| - 1$. In either case, $|G_{l+1}| \geq |G_l| - 2$ and $|B_{l+1}| \leq |B_l| + 1$.

To sum up, in each round l of random priority, we can assume the size of B_l increases by at most 1 and the size of G_l decreases by at most 2. Given this and that $|G_0| = k$ and $|B_0| = n - k$ and that $|G_l| > 0$ for $l \leq \lfloor k/2 \rfloor$, we get

$$\mathbb{E} \left[\sum_{i=1}^n u_i(RP(\mathbf{u})_i) \right] \geq \sum_{l=0}^n \left(\frac{|G_l|}{|G_l| + |B_l|} - \epsilon \right) \geq \sum_{l=0}^{\lfloor \frac{k}{2} \rfloor} \frac{k - 2l}{n - l} - n\epsilon$$

and the ratio is

$$\begin{aligned} \frac{\mathbb{E} [\sum_{i=1}^n u_i(RP(\mathbf{u})_i)]}{w^*(\mathbf{u})} &\geq \frac{\sum_{l=0}^{\lfloor \frac{k}{2} \rfloor} \frac{k - 2l}{n - l} - n\epsilon}{k + \frac{1}{n^2}} \geq \frac{\sum_{l=0}^{\lfloor \frac{k}{2} \rfloor} \frac{k - 2l}{n - l} - n\epsilon}{2k} \\ &= \sum_{l=0}^{\lfloor \frac{k}{2} \rfloor} \frac{1 - \frac{2l}{k}}{2(n - l)} - \frac{n\epsilon}{2k} > \sum_{l=0}^{\lfloor \frac{k}{2} \rfloor} \frac{1 - \frac{2l}{k}}{2n} - \frac{n\epsilon}{2k} \geq \frac{k - 11}{8n} - \frac{n\epsilon}{2k}. \end{aligned}$$

The bound is clearly minimum when k is minimum, that is, $k = \sqrt{n}$. Since this bound holds for any $\mathbf{u} \in C_\epsilon^n$, we get

$$ar(RP) = \inf_{\mathbf{u} \in C_\epsilon^n} \frac{\mathbb{E} [\sum_{i=1}^n u_i(J(\mathbf{u})_i)]}{w^*(\mathbf{u})} \geq \frac{\sqrt{n} - 11}{8n} - \frac{n\epsilon}{2\sqrt{n}}.$$

We can choose ϵ so that the approximation ratio is at least $\frac{1}{20\sqrt{n}}$ for $n \geq 400$ and for $n \leq 400$, the bound holds trivially since random priority matches at least one agent with its favorite item. \square

Next, we state the following lemma about ordinal mechanisms. We leave the details for the full version.

Lemma 5. *Let J be any ordinal mechanism for unit-range representation. Then $ar(J) = O(n^{-1/2})$.*

Our final lemma provides a matching upper bound on the approximation ratio of any truthful-in-expectation mechanism.

Lemma 6. *Let J be a truthful-in-expectation mechanism for unit-range representation. Then $ar(J) = O(n^{-1/2})$.*

Proof. By Lemma 1, we can assume that Mechanism J is anonymous. Let $k \geq 1$ be a parameter to be chosen later and let $\mathbf{u} = (u_1, u_2, \dots, u_n)$ be the valuation profile where

$$u_i(j) = \begin{cases} 1, & \text{for } j = i \\ \frac{2}{k} - \frac{j}{n}, & \text{for } 1 \leq j \leq k+1, j \neq i \\ \frac{n-j}{n^2}, & \text{otherwise} \end{cases} \quad \forall i \in \{1, \dots, k+1\},$$

$$u_i(j) = \begin{cases} 1, & \text{for } j = 1 \\ \frac{2}{k} - \frac{j}{n}, & \text{for } 2 \leq j \leq k+1 \\ \frac{n-j}{n^2}, & \text{otherwise} \end{cases} \quad \forall i \in \{k+2, \dots, n\}.$$

For $i = 2, \dots, k+1$, let $\mathbf{u}^i = (u'_i, u_{-i})$ be the valuation profile where all agents besides agent i have the same valuations as in \mathbf{u} and $u'_i = u_{k+2}$. Note that when agent i on valuation profile \mathbf{u}^i , reports u_i instead of u'_i , the resulting valuation profile is \mathbf{u} . Since J is anonymous and $u'_i = u_1 = u_{k+2} = \dots = u_n$, then agent i receives at most a uniform lottery among these agents on valuation profile \mathbf{u}^i and so it holds that

$$\begin{aligned} \mathbb{E}[u'_i(J(\mathbf{u}^i)_i)] &\leq \frac{1}{n-k+1} + \sum_{j=2}^{k+1} \frac{1}{n-k+1} \left(\frac{2}{k} - \frac{j}{n} \right) + \sum_{j=k+2}^n \frac{1}{n-k+1} \cdot \frac{n-j}{n^2} \\ &\leq \frac{4}{n-k+1}. \end{aligned}$$

Next observe that since J is truthful-in-expectation, agent i should not increase its expected utility by misreporting u_i instead of u'_i on valuation profile \mathbf{u}^i , that is,

$$\mathbb{E}[u'_i(J(\mathbf{u}^i)_i)] \geq \mathbb{E}[u'_i(J(\mathbf{u})_i)]. \quad (1)$$

For all $i = 2, \dots, k+1$, let p_i be the probability that $J(\mathbf{u})_i = i$. Then,

$$\mathbb{E}[u'_i(J(\mathbf{u})_i)] \geq p_i \left(\frac{2}{k} - \frac{i}{n} \right) \geq p_i \left(\frac{2}{k} - \frac{k+1}{n} \right),$$

and by Inequality (1) we get

$$\begin{aligned} p_i \left(\frac{2}{k} - \frac{k+1}{n} \right) &\leq \frac{4}{n-k+1} \\ \Rightarrow p_i &\leq \frac{4}{n-k+1} \cdot \frac{kn}{2n-k(k+1)} \leq \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}. \end{aligned}$$

Let $p = \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}$, i.e. for all i , $p_i \leq p$. We will next calculate an upper bound on the expected social welfare achieved by J on valuation profile \mathbf{u} .

For item $j = 1$, the contribution to the social welfare is upper bounded by 1. Similarly, for each item $j = k+2, \dots, n$, its contribution to the social welfare is upper bounded by $1/n$. Overall, the total contribution by item 1 and items $k+2, \dots, n$ will be upper bounded by 2.

We next consider the contribution to the social welfare from items $j = 2, \dots, k+1$. Define the random variables

$$X_j = \begin{cases} 1, & \text{if } J(\mathbf{u})_j = j \\ \frac{2}{k} - \frac{j}{n}, & \text{otherwise} \end{cases}.$$

The contribution from items $j = 2, \dots, k+1$ is then $\sum_{j=2}^{k+1} X_j$ and so we get

$$\mathbb{E} \left[\sum_{j=2}^{k+1} X_j \right] = \sum_{j=2}^{k+1} \mathbb{E}[X_j] \leq \sum_{j=2}^{k+1} \left(p + \frac{2}{k} - \frac{j}{n} \right) \leq kp + 2.$$

Overall, the expected social welfare of mechanism J is at most $4 + pk$ while the social welfare of the optimal matching is $k+1 + \sum_{i=k+2}^n \frac{n-i}{n^2}$ which is at least k . Since $p = \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}$, the approximation ratio of J then is

$$ar(J) \leq \frac{4 + pk}{k} = \frac{4}{k} + \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}.$$

Let $k = \lfloor \sqrt{n} \rfloor - 1$ and note that $\sqrt{n} - 2 \leq k \leq \sqrt{n} - 1$. Then,

$$\begin{aligned} ar(J) &\leq \frac{4}{k} + \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2} \leq \frac{4}{\sqrt{n}-2} + \frac{4}{n-\sqrt{n}+1} \cdot \frac{(\sqrt{n}-1)n}{2n-(\sqrt{n})^2} \\ &\leq \frac{4}{\sqrt{n}-2} + \frac{4}{\sqrt{n}} \leq \frac{12}{\sqrt{n}} + \frac{4}{\sqrt{n}} = \frac{16}{\sqrt{n}}. \end{aligned}$$

The last inequality holds for $n \geq 9$ and for $n < 9$ the bound holds trivially. This completes the proof. \square

4 Unit-sum valuation functions

In this section, we assume that the representation of the valuation functions is unit-sum. We prove Theorem 1 using the following three lemmas. The first lemma provides a lower bound on the approximation ratio of random priority.

Lemma 7. *For unit-sum representation, $ar(RP) = \Omega(n^{-1/2})$.*

Proof. Let \mathbf{u} be any unit-sum valuation profile and let C be the constant in the bound from Lemma 4. Suppose first that the $w^*(\mathbf{u}) < 4\sqrt{n}/C$. We will show that random priority guarantees an expected social welfare of 1, which proves the lower bound for this case. Consider any agent i and notice that in random priority, the probability that the agent is picked by the l 'th round is l/n , for any $1 \leq l \leq n$ and hence the probability of the agent getting one of its l most preferred items is at least l/n . Let u_i^l be agent i 's valuation for its l 'th most preferred item; agent i 's expected utility for the first round is then at least u_i^1/n . For the second round, in the worst case, agent i 's most preferred item has already been matched to a different agent and so the expected utility of the agent for the first two rounds is at least $u_i^1/n + u_i^2/n$. By the same argument, agent i 's expected utility after n rounds is at least $\sum_{l=1}^n u_i^l/n = 1/n$. Since this holds for each of the n agents, the expected social welfare is at least 1.

Suppose now $w^*(\mathbf{u}) \geq 4\sqrt{n}/C$. We will transform \mathbf{u} to a unit-range valuation profile \mathbf{u}'' . By Lemma 2, the optimal allocation can be achieved by a run of random priority, so we know that in the optimal allocation at most 1 agent will be matched with its least preferred item. Now consider the valuation profile \mathbf{u}' where each agent i 's valuation for its least preferred item is set to 0 (unless it already is 0) and the rest of the valuations are as in \mathbf{u} . Since the ordinal preferences of agents are unchanged, random priority performs worse on this valuation profile, and because of Lemma 2, $w^*(\mathbf{u}') \geq w^*(\mathbf{u}) - 1/n$. Next consider the valuation profile

$$\mathbf{u}'' = \begin{pmatrix} \mathbf{u}' & \mathbf{1} \\ \mathbf{o}^T & 1 \end{pmatrix}.$$

where $\mathbf{o} \in \mathbb{R}^n$ and $\mathbf{o}_j = (j-1)/n^5$. That is, \mathbf{u}'' has $n+1$ agents and items, where agents $1, \dots, n$ have the same valuations for items $1, \dots, n$ as in \mathbf{u}' , every agent has a valuation of 1 for item $n+1$, and agent $n+1$ only has a significant valuation for item $n+1$. Notice that \mathbf{u}'' is a unit-range valuation profile, and $w^*(\mathbf{u}'') \geq w^*(\mathbf{u}') + 1$. Furthermore, $\mathbb{E}[\sum_{i=1}^n u_i(RP(\mathbf{u}'))] \geq \mathbb{E}[\sum_{i=1}^n u_i(RP(\mathbf{u}''))] - 2$ and hence

$$\begin{aligned} \frac{\mathbb{E}[\sum_{i=1}^n u_i(RP(\mathbf{u}))]}{w^*(\mathbf{u})} &\geq \frac{\mathbb{E}[\sum_{i=1}^n u_i(RP(\mathbf{u}'))]}{w^*(\mathbf{u}') + 1/n} \geq \frac{\mathbb{E}[\sum_{i=1}^n u_i(RP(\mathbf{u}''))] - 2}{w^*(\mathbf{u}'') + 1/n - 1} \\ &\geq \frac{\mathbb{E}[\sum_{i=1}^n u_i(RP(\mathbf{u}''))]}{w^*(\mathbf{u}'')} - \frac{2}{w^*(\mathbf{u}'')} \geq \frac{C}{\sqrt{n}} - \frac{2}{w^*(\mathbf{u})} \\ &\geq \frac{C}{\sqrt{n}} - \frac{2}{4\sqrt{n}/C} = \frac{C}{2\sqrt{n}}. \end{aligned}$$

This completes the proof. \square

Next, we bound the approximation ratio of any ordinal (not necessarily truthful-in-expectation) mechanism. The proof is similar to that of Lemma 5.

Lemma 8. *Let J be an ordinal mechanism for unit-sum representation. Then $ar(J) = O(n^{-1/2})$.*

Finally, the upper bound for any truthful-in-expectation mechanism is given by the following lemma.

Lemma 9. *Let J be a truthful-in-expectation mechanism for unit-sum representation. Then $ar(J) = O(n^{-1/2})$.*

For the proof, the main observation is that the valuation profile used in the proof of Lemma 6 can easily be modified in a way such that all rows of the matrices of valuations sum up to one, to obtain a unit-sum valuation profile. Then exactly the same steps used in the proof of Lemma 6 also prove Lemma 9.

5 Extensions

Allowing ties

Our results extend if we allow ties in the image of the valuation function. All of our upper bounds hold trivially. For the approximation guarantee of random priority, first the mechanism clearly must be equipped with some tie-breaking rule to settle cases where indifferences appear. For all natural (fixed before the execution of the mechanism) tie-breaking rules the lower bounds still hold. To see this, consider any valuation profile with ties and a tie-breaking rule for random priority. We can simply add sufficiently small quantities ϵ_{ij} to the valuation profile according to the tie-breaking rule and create a new profile without ties. The assignment probabilities of random priority will be exactly the same as for the version with ties, and random priority achieves an $\Omega(1/\sqrt{n})$ approximation ratio on the new profile. Then since ϵ_{ij} were sufficiently small, the same bound holds for the original valuation profile.

$[0, 1]$ valuation functions

All of our results apply to the extension of the unit-range representation where 0 is not required to be in the image of the function, that is $\max_j u_i(j) = 1$ and for all j , $u_i(j) \in [0, 1]$. This representation captures scenarios where agents can be more or less indifferent between every single item. Since every unit-range valuation profile is also a valid profile for this representation, the upper bounds hold trivially. For the approximation ratio of random priority, we obtain the following corollary.

Corollary 1. *The approximation ratio of random priority, for the setting with $[0, 1]$ valuation functions is $\Omega(n^{-1/2})$.*

The intuition for the proof is that, similarly to the proof of Lemma 7, we can “turn” a $[0, 1]$ profile into a unit-range profile with only a constant change in the approximation ratio and then use our bound for unit-range profiles. The details are left for the full version.

References

1. Atila Abdulkadiroglu and Tayfun Sönmez. Matching Markets: Theory and Practice. *Advances in Economics and Econometrics (Tenth World Congress)*, pages 3–47, 2013.
2. Marek Adamczyk, Sankowski Piotr, and Qiang Zhang. Efficiency of Truthful and Symmetric Mechanisms in One-sided Matching. In *Proceedings of the 7th Symposium of Algorithmic Game Theory. To appear*, 2014.
3. Elliot Anshelevich and Sanmay Das. Matching, cardinal utility, and social welfare. *ACM SIGECom Exchanges*, 9(1):4, 2010.
4. Salvador Barbera. Strategy-proof Social Choice. In K. J. Arrow, A. K. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 2, chapter 25. North-Holland: Amsterdam, 2010.
5. Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. Social Welfare in One-Sided Matching Markets without Money. In *APPROX-RANDOM*, pages 87–98, 2011.
6. Anna Bogomolnaia and Hervé Moulin. A New Solution to the Random Assignment Problem. *Journal of Economic Theory*, 100:295–328, 2001.
7. Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. Optimal social choice functions: a utilitarian view. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 197–214. ACM, 2012.
8. Deeparnab Chakrabarty and Chaitanya Swamy. Welfare maximization and truthfulness in mechanism design with ordinal preferences. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 105–120, 2014.
9. Shaddin Dughmi and Arpita Ghosh. Truthful assignment without money. In *ACM Conference on Electronic Commerce*, pages 325–334, 2010.
10. Uriel Feige and Moshe Tennenholtz. Responsive lotteries. In *Algorithmic Game Theory*, pages 150–161. Springer, 2010.
11. Aris Filos-Ratsikas and Peter Bro Miltersen. Truthful approximations to range voting. *CoRR*, abs/1307.1766, 2013.
12. Mingyu Guo and Vincent Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, pages 881–888, 2010.
13. Aanund Hylland and Richard Zeckhauser. The Efficient Allocation of Individuals to Positions. *The Journal of Political Economy*, 87(2):293–314, 1979.
14. Timo Mennle and Sven Seuken. An axiomatic approach to characterizing and relaxing strategyproofness of one-sided matching mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 37–38, 2014.
15. Ariel D. Procaccia. Can Approximation Circumvent Gibbard-Satterthwaite? In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*. AAAI Press, 2010.
16. Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 177–186. ACM, 2009.
17. Tayfun Sönmez and Utku Ünver. Matching, allocation and exchange of discrete resources. *Handbook of Social Economics*, 1A:781–852, 2011.
18. Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.
19. Lin Zhou. On a Conjecture by Gale about One-Sided Matching Problems. *Journal of Economic Theory*, 52:123–135, 1990.