University of Pennsylvania
ScholarlyCommons

Departmental Papers (CIS) | Department of Computer & Information Science

8-1-2011

Social Welfare in One-Sided Matching Markets Without Money

Anand Bhalgat
University of Pennsylvania

Deeparnab Chakrabarty
University of Pennsylvania

Sanjeev Khanna
University of Pennsylvania, sanjeev@cis.upenn.edu

Bhalgat, A., Chakrabarty, D., & Khanna, S., Social Welfare in One-Sided Matching Markets Without Money, *Approximation, Randomization, and Combinatorial Optimization, Algorithms & Techniques, 14th International Workshop, APPROX 2011, and 15th International Workshop, RANDOM 2011*, Aug. 2011, doi: 10.1007/978-3-642-22935-0_8
Copyright © 2011, Springer Berlin / Heidelberg

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/cis_papers/663
For more information, please contact repository@pobox.upenn.edu.

Social Welfare in One-Sided Matching Markets Without Money

Abstract
We study social welfare in one-sided matching markets where the goal is to efficiently allocate n items to n agents that each have a complete, private preference list and a unit demand over the items. Our focus is on allocation mechanisms that do not involve any monetary payments.We consider two natural measures of social welfare: the ordinal welfare factor which measures the number of agents that are at least as happy as in some unknown, arbitrary benchmark allocation, and the linear welfare factor which assumes an agent's utility linearly decreases down his preference lists, and measures the total utility to that achieved by an optimal allocation. We analyze two matching mechanisms which have been extensively studied by economists. The first mechanism is the random serial dictatorship (RSD) where agents are ordered in accordance with a randomly chosen permutation, and are successively allocated their best choice among the unallocated items. The second mechanism is the probabilistic serial (PS) mechanism of Bogomolnaia and Moulin [8], which computes a fractional allocation that can be expressed as a convex combination of integral allocations. The welfare factor of a mechanism is the infimum over all instances. For RSD, we show that the ordinal welfare factor is asymptotically 1/2, while the linear welfare factor lies in the interval [.526, 2/3]. For PS, we show that the ordinal welfare factor is also 1/2 while the linear welfare factor is roughly 2/3. To our knowledge, these results are the first non-trivial performance guarantees for these natural mechanisms.

Disciplines
Computer Sciences

Comments
Bhalgat, A., Chakrabarty, D., & Khanna, S., Social Welfare in One-Sided Matching Markets Without Money, *Approximation, Randomization, and Combinatorial Optimization, Algorithms & Techniques, 14th International Workshop, APPROX 2011, and 15th International Workshop, RANDOM 2011*, Aug. 2011, doi: 10.1007/978-3-642-22935-0_8

Copyright © 2011, Springer Berlin / Heidelberg

This conference paper is available at ScholarlyCommons: http://repository.upenn.edu/cis_papers/663

Social Welfare in One-Sided Matching Markets without Money

Anand Bhalgat*, Deeparnab Chakrabarty, and Sanjeev Khanna**

Department of Computer and Information Science,
University of Pennsylvania
bhalgat@seas.upenn.edu, deepc@seas.upenn.edu, sanjeev@cis.upenn.edu

Abstract. We study social welfare in one-sided matching markets where the goal is to efficiently allocate n items to n agents that each have a complete, private preference list and a unit demand over the items. Our focus is on allocation mechanisms that do not involve any monetary payments. We consider two natural measures of social welfare: the *ordinal welfare factor* which measures the number of agents that are at least as happy as in some unknown, arbitrary benchmark allocation, and the *linear welfare factor* which assumes an agent's utility linearly decreases down his preference lists, and measures the total utility to that achieved by an optimal allocation.

We analyze two matching mechanisms which have been extensively studied by economists. The first mechanism is the random serial dictatorship (RSD) where agents are ordered in accordance with a randomly chosen permutation, and are successively allocated their best choice among the unallocated items. The second mechanism is the probabilistic serial (PS) mechanism of Bogomolnaia and Moulin [8], which computes a fractional allocation that can be expressed as a convex combination of integral allocations. The welfare factor of a mechanism is the infimum over all instances. For RSD, we show that the ordinal welfare factor is asymptotically 1/2, while the linear welfare factor lies in the interval $[.526, 2/3]$. For PS, we show that the ordinal welfare factor is also 1/2 while the linear welfare factor is roughly 2/3. To our knowledge, these results are the first non-trivial performance guarantees for these natural mechanisms.

1 Introduction

In the one-sided matching market problem[1], the goal is to efficiently allocate n items, I, to n unit-demand agents, A, with each agent a having a complete and private preference list $\geq_a$ over these items. The problem arises in various applications such as assigning dormitory rooms to students, time slots to users of a common machine, organ allocation markets, and so on. Since the preferences

* Supported in part by NSF Awards CCF-0635084 and IIS-0904314.

[1] In the literature, the problem has been alternately called the *house allocation* or *assignment* problem.

L.A. Goldberg et al. (Eds.): APPROX/RANDOM 2011, LNCS 6845, pp. 87–98, 2011.
© Springer-Verlag Berlin Heidelberg 2011

are private, we focus on *truthful* (strategyproof) mechanisms in which agents do not have an incentive to misrepresent their preferences. One class of such mechanisms involve monetary compensations/payments among agents. However, in many cases (e.g., in the examples cited above), monetary transfer may be infeasible due to reasons varying from legal restrictions to plain inconvenience. Hence, we focus on truthful mechanisms without money.

A simple mechanism for the one-sided matching problem is the following: agents arrive one-by-one according to a fixed order σ picking up their most preferred unallocated item. This is called as a *serial dictatorship* mechanism. The *random serial dictatorship* (RSD) mechanism picks the order σ uniformly at random among all permutations. Apart from being simple and easy to implement, RSD has attractive properties: it is truthful, fair, anonymous/neutral, non-bossy[2], and returns a Pareto optimal allocation. In fact, it is the only truthful mechanism with the above properties [26], and there is a large body of economic literature on this mechanism (see Section 1.2).

Despite this, an important question has been left unaddressed: how *efficient* is this mechanism? To be precise, what is the guarantee one can give on the social welfare obtained by this algorithm when compared to the optimal social welfare? As computer scientists, we find this a natural and important question, and we address it in this paper.

The usual recourse to measure the social welfare of a mechanism is to *assume* the existence of *cardinal* utilities u_{ij} of agent i for item j with the semantic that agent i prefers item j to ℓ iff $u_{ij} > u_{i\ell}$. A mechanism has *welfare factor* α if for every instance the utility of the matching returned is at least α times that of the optimum utility matching. There are a couple of issues with this. Firstly, nothing meaningful can be said about the performance of RSD if the utilities are allowed to be arbitrary. This is because the optimum utility matching might be arising due to one particular agent getting one particular item (a single edge), however with high probability, any random permutation would lead to another agent getting the item and lowering the total welfare by a lot [3]. Secondly, the assumption of cardinal utilities inherently ties up the performance of the algorithm with the 'cardinal numbers' involved; the very quantities whose existence is only an assumption. Rather, what is needed is an *ordinal* scale of analyzing the quality of a mechanism; a measure that depends only on the order/preference lists of the agents rather than the precise utility values.

In this paper, we propose such a measure which we call the *ordinal social welfare* of a mechanism. Given an instance of items and agents with their preference lists, we assume that there exists some benchmark matching M^*, unknown to the mechanism. We stress here this can be *any* matching. We say that the *ordinal welfare factor* of a (randomized) mechanism is α, if for any instance and

[2] A mechanism is neutral if the allocation of items doesn't change with renaming, and is non-bossy if no agent can change his preference so that his allocation remains unchanged while someone else's changes.

[3] The reader may notice similarities of RSD with online algorithms for bipartite matching problems. We elaborate on the connection in Section 2.2.

every matching M^*, at least αn agents (in expectation) get an item which they prefer at least as much as what they get in M^*.

A discussion of this measure is in order. Firstly, the measure is *ordinal* and is well defined whenever the utilities are expresses via preference lists. Secondly, the notion is independent of any 'objective function' that an application might give rise to since it measures the ordinal social welfare with respect to any desired matching. One disadvantage of the concept is that it is global: it counts the fraction of the *total* population which gets better than their optimal match. In other words, if everyone is 'happy' in the benchmark matching M^*, then a mechanism with the ordinal welfare factor α will make an α fraction of the agents happy. However if M^* itself is *inefficient*, say only 1% of the agents are 'happy' in M^*, then the ordinal welfare factor does not say much. For instance, it does not help for measures like "maximize number of agents getting their first choice", for in some instances, this number could be tiny in any M^*. Furthermore, it does not say anything about the "fairness" of the mechanism, e.g. a mechanism may have the ordinal welfare factor close to 1, but there may exists an agent who is almost always allocated an item that he prefers less than M^*. Finally, we observe that the ordinal welfare factor of any mechanism, even ones which know the true preference lists, cannot be larger than 1/2. The reason for this is that the allocation must be competitive with respect to all benchmark matchings simultaneously, and it can be seen (Theorem 8) that in the instance when all agents have the same preference list, if M^* is chosen to be a random allocation, then no mechanism can have an ordinal welfare factor better than 1/2. Our first result is that the ordinal welfare factor of RSD is in fact asymptotically 1/2.

Theorem 1. *The ordinal welfare factor of RSD is at least* $1/2 - o(1)$.

Till now we have focussed on the RSD mechanism since it is a simple (and essentially unique) truthful mechanism for the matching market problem. A mechanism is called truthful if misrepresenting his preference list doesn't strictly increase the total *utility* of an agent, where the utility is defined as the cardinal utility obtained by the agent on getting his allocated item. However, when the utilities of agents are represented as preference lists, one needs a different definition. In light of this, Bogomolnaia and Moulin [8] proposed a notion of truthfulness based on the *stochastic dominance*: for an agent a random allocation rule stochastically dominates another if the probability of getting one of his top k choices in the first rule is at least that in the second, for any k. A mechanism is called (weakly) truthful if no agent can obtain a stochastically dominating allocation by misreporting his preference list. With this definition, the authors propose a mechanism called the *probabilistic serial* (PS) algorithm, and prove that it is weakly truthful; the mechanism is illustrated in Section 1.1.

PS and RSD are incomparable and results on RSD do not a priori imply those for PS, nevertheless, PS has an ordinal welfare factor of 1/2 as well.

Theorem 2. *The ordinal welfare factor of PS algorithm is at least* 1/2.

Ordinal Welfare Factor and Popular Matchings Our notion of ordinal welfare factor is somewhat related to the notion of *popular* matchings [14,3,21]. Given

preference lists of agents, a matching M is said to be more popular than M' if the number of agents getting *strictly* better items in M is at least the number of agents getting strictly better items in M'. A matching is popular if no other matching is more popular than it. Thus while comparing a matching M to M', the notion of popular matchings distinguishes between agents that prefer M and agents that are neutral, unlike in the case of ordinal welfare factor.

It can be easily seen that any popular matching has an ordinal welfare factor of at least 1/2, however, (a) not every input instance has a popular matching, and (b) no truthful algorithms are known to compute them when they exist. A few modified measures such as unpopularity factor, unpopularity margin and popular mixed matching have also been studied in the literature [22,17,21].

Linear Utilities. We also analyze the performance of RSD and PS mechanisms when agents' utilities are linear - arguably, one of the most commonly studied special case of cardinal utilities. In this model, we assume that the utility for an agent for his i^{th} preference is $\frac{n-i+1}{n}$. Observe that *any* serial dictatorship mechanism achieves a welfare of at least $(n+1)/2$ since the agent at step t gets his t^{th} choice or better, giving him a utility of at least $(1-(t-1)/n)$. How much better does RSD do? Intuitively, one would expect the worst case instance would be one where each agent gets one of his top $o(n)$ choices; that would make the optimum value $n-o(n)$. We call such instances as *efficient* instances since there is an optimum matching where every one gets their (almost) best choice. We show that for efficient instances, RSD's utility is at least $\frac{2n}{3}-o(n)$, and there exists instances where RSD does no better. These bounds hold for PS as well.

Theorem 3. *With linear utilities and efficient instances, RSD has linear welfare at least* $2/3-o(1)$*, and there exist efficient instances for which this is tight.*

Theorem 4. *With linear utilities and efficient instances, PS has linear welfare at least* $2/3-o(1)$*, and there exist efficient instances for which this is tight.*

The following theorem summarizes our results on general instances, and we refer the reader to the full version of this paper [6] for its proof.

Theorem 5. *On general instances, the linear welfare factors of RSD and PS algorithms are at least* 0.526 *and* 0.6602 *respectively.*

Extensions. We consider two extensions to our model and focus on the performance of RSD, leaving that of PS as an open direction. In the first, we let the preference lists be incomplete. The proof of Theorem 1 implies that the ordinal welfare factor of RSD remains unchanged. For linear utilities, we generalize the definition as follows: for an agent with a preference list of length ℓ, the ith choice gives him a utility of $(\ell-i+1)/\ell$. We show that RSD doesn't perform very well.

Theorem 6. *For linear utilities, RSD gets at least* $\tilde{\Omega}(n^{-1/3})$ *fraction of the social optimum. Furthermore, there are instances, where the welfare of RSD is at most* $\tilde{O}(n^{-1/3})$ *fraction of the social optimum.*

In the second extension, we let the demand of an agent be for sets of size K or less, for some $K \geq 1$. Agents now arrive and pick their best 'bundle' among the unallocated items. The ordinal welfare factor of a mechanism is now α if at least an α fraction of agents get a bundle that is as good (assuming there is a complete order on the set of bundles) as what they got in an arbitrary benchmark allocation. We show that RSD has ordinal welfare factor $\Theta(1/K)$.

Theorem 7. *In the case when each agent has a maximum demand of K items, the ordinal welfare factor of RSD is $\Theta(1/K)$.*

1.1 Preliminaries

Utility Models, Truthful Mechanisms, Welfare Factors. As stated above, we consider two models for utilities of agents. In the *cardinal utility model*, each agent a has a utility function $u_a : I \to \mathbb{R}_{\geq 0}$, with the property that $j >_a \ell$ iff $u_a(j) > u_a(\ell)$. Given a distribution on the matchings, the utility of agent a is $u_a(M) := \sum_{M \in \mathcal{M}} p(M) u_a(M(a))$, where $p(M)$ is the probability of matching M. In this paper, we focus on the special case of *linear utility model* where the ith ranked item for any agent a gives him a utility of $(1-(i-1)/n)$. We call an instance *efficient*, if there is a matching which matched every agent to an item in his top $o(n)$ (for concreteness, let's say this is $n^{1/5}$) choices. In the *ordinal utility model*, each agent a represents his utility only via his complete preference list $\geq_a$ over the items. A mechanism $\mathcal{A}$ is *truthful* if no agent can misrepresent his preference and obtain a better item. In the cardinal utility model this implies that for all agents a and utility functions u_a, u'_a, we have $u_a(M) \geq u_a(M')$ where $M = \mathcal{A}(u_1, \ldots, u_n)$ and $M' = \mathcal{A}(u_1, \ldots, u'_a, \ldots, u_n)$. A mechanism has *linear welfare factor* of α if for all instances the (expected) sum of linear utilities of agents obtained from the allocation of the mechanism is at least α times the optimal utility allocation for that instance. A mechanism has *ordinal welfare factor* of α if for all instances, and for *all matchings* M^*, at least α fraction of agents (in expectation) get an item at least as good as that allocated in M^*.

The Probabilistic Serial Mechanism. The *probabilistic serial* (PS) mechanism was suggested by Bogomolnaia and Moulin [8]. The mechanism fractionally allocates items to agents over multiple phases, we denote the fraction of the item i allocated to an agent a by $x(a,i)$. These fractions are such that $\sum_{a \in A} x(a,i) = \sum_{i \in I} x(a,i) = 1$ for all agents a and items i. Thus this fractional allocation defines a distribution on integral matchings.Initially, $x(a,i) = 0$ for every agent a and item i. We say that an item i is allocated if $\sum_{a \in A} x(a,i) = 1$, otherwise we call it to be available. The algorithm grows $x(a,i)$'s in phases, and in each phase one or more items get completely allocated. During a phase of the algorithm, each agent a grows $x(a,i)$ at the rate of 1 where i is his best choice in the set of available items. The current phase completes and the new phase starts when at least one item that was available in the current phase, gets completely allocated. The algorithm continues until all items are allocated.

We make a few observations about the above algorithm which will be useful in our analysis: (a) the algorithm terminates at time $t = 1$, at which time all agents

are fractionally allocated one item, that is, $\sum_{i \in I} x(a,i) = 1$, (b) any phase lasts for time at least $1/n$ and at most 1, and (c) by time $< j/n$ for any $1 \leq j \leq n$, at most $(j-1)$ phases are complete.

1.2 Other Related Work

There is a huge amount of literature on matching markets starting with the seminal paper of Gale and Shapley [13], see [24,25,2] for detailed surveys. The one-sided matching market design problem was first studied by Hylland and Zeckhauser [18] who propose a mechanism to find a distribution on matchings via a market mechanism. Their mechanism returns Pareto optimal, envy-free solutions, but is not truthful. Zhou [27], showed that there can be no truthful mechanism which is anonymous/neutral and satisfies *ex ante* Pareto optimality. Svensson [26] showed that serial dictatorship mechanisms are the only truthful mechanisms which are (ex post) Pareto optimal, non bossy, and anonymous.

The study of mechanisms with ordinal utilities for this problem was started by Bogomolnaia and Moulin[8]. The PS mechanism was proposed in an earlier paper by Cres and Moulin [11]. Following the work of [8], there was a list of work characterizing stochastic dominance [1,9], and generalizing it to the case of incomplete preference lists [20], and to multiple copies of items [10]. The study of mechanism design without money has also been of recent interest in the computer science community, see e.g. [23,5,12,16].

2 Ordinal Welfare Factor of RSD and PS Mechanisms

In this section, we prove Theorems 1 and 2. We first show that the ordinal welfare factor of *any* mechanism is at most 1/2 in the instance where every agent has the same preference list.

Theorem 8. *If every agent has the same preference list* $(1, 2, \ldots, n)$*, then the ordinal welfare factor of any mechanism is at most* $1/2 + 1/2n$*.*

Proof. A mechanism returns a probability distribution on matchings which we will interpret as a distribution of permutations. Let $\mathcal{D}$ be that distribution. We choose the benchmark matching M^* to be a random perfect matching. It suffices to show that for any fixed permutation $\sigma \in \mathcal{D}$, the expected number of agents a such that $\sigma(a) \leq \pi(a)$ is $(n+1)/2$. Since π is chosen uniformly at random, the probability that $\pi(a) < \sigma(a)$ is precisely $(\sigma(a) - 1)/n$, and so the expected number of happy people for the permutation σ is $(n+1)/2$.

2.1 Ordinal Welfare Factor of RSD

In this section, we prove Theorem 1. Let M^* be the unknown benchmark matching. We call an agent a *dead* at time t if he hasn't arrived yet and all items as good as $M^*(a)$ in his preference list has been allocated. Let D_t be the expected number of dead agents at time t. Let $\texttt{ALG}_t$ be the expected number of agents who

get an item as good as their choice in M^* by time t. From the above definition, we get

$$\texttt{ALG}_{t+1} - \texttt{ALG}_t = 1 - \frac{D_t}{n-t} \tag{1}$$

We will now bound D_t from above which along with (1) will prove the theorem.

Lemma 1. *$D_t \leq \frac{(t+2)(n-t)}{n+1}$ for $1 \leq t \leq n$.*

Before proving the lemma, note that adding (1) for $t = 1$ to $n-1$ gives $\texttt{ALG}_n - \texttt{ALG}_1 \geq \sum_{t=1}^{n-1} \left(1 - \frac{t+2}{n+1}\right)$, implying $\texttt{ALG}_n - \texttt{ALG}_1 \geq n/2 - 2n/n$. This proves that the ordinal welfare factor of RSD is at least $1/2 - o(1)$ proving Theorem 1.

Proof. Let us start with a few definitions. For an item i and time t, let $\texttt{ALL}_{i,t}$ be the event that item i is allocated by time t. For an agent a and time t, let $\texttt{LATE}_{a,t}$ be the event that a arrives after time t. The first observation is this: if an agent a is dead at time t, then the event $\texttt{ALL}_{M(a),t}$ and $\texttt{LATE}_{a,t}$ must have occurred. Therefore we get

$$D_t \leq \sum_{a \in A} \Pr[\texttt{ALL}_{M(a),t} \ \wedge \ \texttt{LATE}_{a,t}] \tag{2}$$

Note that $\Pr[\texttt{LATE}_{a,t}]$ is precisely $(1 - t/n)$. Also, note that $\sum_{i \in I} \Pr[\texttt{ALL}_{i,t}] = t$. This is because all agents are allocated *some* item. Now suppose *incorrectly* that $\texttt{ALL}_{M(a),t}$ and $\texttt{LATE}_{a,t}$ were independent. Then, (2) would give us

$$D_t \leq (1 - \frac{t}{n}) \sum_{a \in A} \Pr[\texttt{ALL}_{M(a),t}] = (1 - \frac{t}{n}) \sum_{i \in I} \Pr[\texttt{ALL}_{i,t}] = \frac{t(n-t)}{n} \tag{3}$$

which is at most the RHS in the lemma. However, the events are not independent, and one can construct examples where the above bound is indeed incorrect. To get the correct bound, we need the following claim.

Claim.

$$\frac{\Pr[\texttt{ALL}_{M(a),t} \ \wedge \ \texttt{LATE}_{a,t}]}{(n-t)} \leq \frac{\Pr[\texttt{ALL}_{M(a),t+1} \ \wedge \ \overline{\texttt{LATE}_{a,t+1}}]}{(t+1)}$$

Proof. This follows from a simple charging argument. Fix a relative order of all agents other than a and consider the n orders obtained by placing a in the n possible positions. Observe that if the event $\texttt{ALL}_{M(a),t} \wedge \texttt{LATE}_{a,t}$ occurs at all, it occurs exactly $(n-t)$ times when a's position is $t+1$ to n. Furthermore, crucially observe that if the position of a is 1 to $t+1$, the item $M(a)$ will still be allocated. This is because the addition of a only leads to worse choices for agents following him and so if $M(a)$ was allocated before, it is allocated even now. This proves that for every $(n-t)$ occurrences of $\texttt{ALL}_{M(a),t} \wedge \texttt{LATE}_{a,t}$, we have $(t+1)$ occurrences of the event $\texttt{ALL}_{M(a),t+1} \wedge \overline{\texttt{LATE}_{a,t+1}}$. The claim follows as it holds for every fixed relative order of other agents.

Now we can finish the proof of the lemma. From Claim 2.1, we get

$$\frac{t+1}{n-t} \cdot \Pr[\texttt{ALL}_{M(a),t} \wedge \texttt{LATE}_{a,t}] \leq \Pr[\texttt{ALL}_{M(a),t+1}] - \Pr[\texttt{ALL}_{M(a),t+1} \wedge \texttt{LATE}_{a,t+1}]$$

Taking the second term of the RHS to the LHS, adding over all agents, and invoking (2), we get

$$\frac{t+1}{n-t} \cdot D_t + D_{t+1} \leq t+1 \tag{4}$$

Using the fact that $D_{t+1} \geq D_t - 1$ (the number of dead guys cannot decrease by more than 1), and rearranging, proves the lemma.

2.2 RSD and Online Bipartite Matching

In this section, we highlight the relation between RSD and algorithms for online bipartite matching. In fact, the analysis of RSD above can be seen as a generalization of online bipartite matching algorithms.

In the online bipartite matching problem, vertices of one partition (think of them as agents) are fixed while vertices of the other partition (think of them as items) arrive in an adversarial order. Karp, Vazirani and Vazirani [19] gave the following algorithm (KVV) for the problem: fix a *random* ordering of the agents, and when an item arrives give it to the first unmatched agent in this order. They proved[4] that the expected size of the matching obtained is at least $(1-1/e)$ times the optimum matching. The KVV theorem can be 'flipped around' to say the following. Suppose each agent has the preference list which goes down its desired items in the order of entry of items. Then, if agents arrive in a random order and pick their best, unallocated, desired item, in expectation an $(1-1/e)$ fraction of agents are matched. That is, if we run RSD on this instance (with incomplete lists), an $(1-1/e)$ fraction of agents will get an item.

The above result does not a priori imply an analysis of RSD, the reason being that in our problem an agent a, when he arrives, is allocated an item even if that item is *worse* than what he gets in the benchmark matching M^*. This might be bad since the allocated item could be 'good' item for agents to come. In particular, if the order chosen is not random but arbitrary, the performance of the algorithm is quite bad; in contrast, the online matching algorithm still has a competitive ratio of 1/2. Nevertheless, similar techniques prove both the results and our analysis can be tailored to give a proof of the online bipartite matching result (See [6] for details).

2.3 Ordinal Welfare Factor of PS

In this section, we prove Theorem 2. We suggest the reader to refer to the algorithm and its properties as described in Section 1.1. In particular, we will use the following observation.

[4] In 2008, a bug was found in the original extended abstract of [19], but was soon resolved. See [15,7,4] for discussions and resolutions.

Observation 1: By time $< j/n$, for any $1 \leq j \leq n$, at most $(j-1)$ items are completely allocated.

Let M^* be the unknown benchmark matching. For an agent a, let t_a be the time at which the item $M^*(a)$ is completely allocated. Observe that the probability agent a gets an item $M^*(a)$ or better is precisely t_a, since till this time $x(a,i)$ increases for items $i \geq_a M^*(a)$. Summing up all agents, we see that the ordinal welfare factor of the PS mechanism is $\sum_a t_a$. The observation above implies at most $(j-1)$ agents have $t_a < j/n$. So, $\sum_a t_a \geq \sum_{j=1}^{n}(n-j+1)/n \geq n/2+1/2$. This completes the proof of Theorem 2.

3 Linear Welfare Factor of RSD and PS

In this section, we establish bounds on the linear welfare factor of RSD and PS mechanisms. We first prove Theorem 3 in two lemmas. Recall that an instance is called efficient if there exists a matching in which every agent is matched to an item in his top $o(n)$ choices.

Lemma 2. *When the instance is efficient, the linear welfare factor of RSD is at least* $(2/3 - o(1))$.

Proof. The proof follows from Lemma 1. Let U_t denote the expected utility obtained by time t. Consider the agent coming at time $t+1$. If he is not dead already, then he will get a utility of at least $(1-o(1))$ (since the instance is efficient). If he is dead, then he will get a utility of at least $(1-t/n)$. This is because only t items have been allocated and this agent takes an item $(t+1)$th ranked or higher. Therefore,

$$U_{t+1} - U_t \geq \left(1 - \frac{D_t}{n-t}\right) \cdot (1-o(1)) + \frac{D_t}{n-t} \cdot (1-t/n) \geq 1 - o(1) - \frac{t}{n} \cdot \frac{D_t}{n-t}$$

Using Lemma 1, we get $U_{t+1} - U_t \geq 1 - o(1) - \frac{t(t+2)}{n(n+1)}$. Summing over all t, we get that the total utility of RSD is at least $(1-o(1))n - (n/3 + o(n)) = (2/3 - o(1))n$.

The above analysis can be modified via a 'balancing trick' to give a strictly better than 50% guarantee for all instances. We refer the reader to [6] for details.

Lemma 3. *When the utilities are linear, there exists an efficient instance for which RSD gets a utility of at most* $(2/3 + o(1))n$.

Proof. Partition n agents and items into t blocks of size n/t each, where $t = n^{1/5}$. We denote the jth block of agents and items by A_j and I_j respectively, and they number from $\left(\frac{(j-1)n}{t} + 1\right)$ to $\frac{jn}{t}$.

We now illustrate the preference lists of agents. Fix an agent a in block A_j. Let he be the kth agent in the block, where $1 \leq k \leq n/t$, i.e. his agent number is $(j-1)n/t + k$. A random set of t^3 items is picked from each of blocks $I_1, \ldots, I_{j-1}$, and these form the first $(j-1)t^3$ items in his preference list, in increasing order

of item number. The item $(j-1)n/t+k$ is his $((j-1)t^3+1)^{th}$ choice. His remaining choices are the remaining items considered in increasing order. This completes the description of the preference lists of the agents.

Note that if every agent a is assigned the corresponding item with the same number, then each agent gets one of his top t^4 choices, leading to a utility of at least $(1-\frac{t^4}{n}) = 1-o(1)$. So, the instance is indeed efficient. We now show that RSD gets utility at most $2n/3+o(1)$.

Let σ be a random permutation of the agents. We divide σ into t chunks of n/t agents, with the jth chunk, S_j, consisting of agents $\sigma(\frac{(j-1)n}{t}+1)$ to $\sigma(\frac{jn}{t})$. Note that with high probability ($\geq (1-1/t^3)$), we have that for any block A_j and chunk S_i, $|A_j \cap S_i| \in \left[(1-\frac{1}{t^2})\frac{n}{t^2}\ , (1+\frac{1}{t^2})\frac{n}{t^2}\right]$. Since agents prefer items in 'higher' blocks to 'lower' blocks, we claim the following.

Claim. With high probability, at least $(1-\frac{1}{t^3})$ fraction of the items in the first i blocks have been allocated after arrival of first i chunks. (Proof omitted; see [6].)

Now we are ready to analyze RSD. Consider the $(i+1)^{th}$ chunk of agents. With high probability, there are at least $\frac{n}{t^2}(1-\frac{1}{t^2})$ agents from each block $A_1,\ldots,A_i$ in S_{i+1}. Since only in/t^3 items remain from the first i block of items, at least $\frac{in}{t^2}(1-\frac{1}{t^2})-\frac{in}{t^3}$ of these agents must get an item from blocks $(i+1)$ or higher. However, this gives them utility at most $(1-\frac{in/t}{n}) \geq 1-i/t$. That is, the *drop* in their utility to what they get in the optimum is at least i/t. Summing the total drop over all agents and all chunks, we get that the difference between RSD and the optimum is at least

$$\sum_{i=1}^{t} \frac{in}{t^2}(1-\frac{1}{t})\frac{i}{t} = (1-o(1))\frac{n}{t^3}\sum_{i=1}^{n} i^2 = n/3$$

Therefore, the social welfare of RSD is at most $(2/3+o(1))n$.

Linear Welfare Factor of PS Mechanism We establish the lower bound in this abstract, and the upper bound instance, which is similar to that for RSD, can be found in [6]. As in the case of RSD, we focus on efficient instances.

Lemma 4. *For efficient instance, the linear welfare factor of PS* $\geq 2/3-o(1)$.

Proof. Let o_a denote the utility obtained by agent a in the utility optimal matching. Since the instance is efficient, $o_a = 1-o(1)$ for all agents a.

Consider the jth phase of PS, and suppose it lasts for time Δ_j. Observation 1 implies that $\sum_{j\leq \ell} \Delta_j \geq \ell/n$. Furthermore, in phase j, at least $(n-j+1)$ agents obtain utility at a rate higher than their utility in the optimal matching. This is because at most $(j-1)$ items have been allocated. Also, the remaining $(j-1)$ agents are getting utility at a rate at least $(1-(j-1)/n)$ since they are growing their $x(a,i)$ on their j^{th} choice or better. So, the total utility obtained by PS is at least $\sum_{j=1}^{n} \Delta_j \cdot \big((n-j+1)\cdot(1-o(1)) + (j-1)\cdot(1-\frac{j-1}{n})\big)$ which evaluates to $\sum_{j=1}^{n} \Delta_j \left(\frac{n^2-(j-1)^2}{n}\right) - o(n)$

The above summation is smallest if Δ_1 is as small as possible, modulo which, Δ_2 is as small as possible and so on. Given the constraint on Δ_j's, we get that this is at least $\sum_{j=1}^{n} \frac{n^2-(j-1)^2}{n^2} = 2n/3 - o(n)$.

4 Concluding Remarks

We first give very brief sketches of the proofs of Theorems 6 and 7. Full proofs can be found in [6].

Incomplete Preference Lists. The ordinal welfare factor of RSD remains the same, however, the linear welfare factor of RSD drops to $\tilde{\Theta}(1/n^{1/3})$. This is because some agents can have 'long' preference lists and some agents have 'short' preference lists, and in a random order the long preference list agents can take away items of the short preference list ones. However, if the lengths of the preference lists of the 'long agents' are 'too long', they get an item with high enough linear utility. The correct balancing argument gives the $\tilde{\Theta}(\frac{1}{n^{1/3}})$ factor.

Non-unit demands. Note that a single agent's choice can disrupt the choices of K other agents. Therefore, it is not too difficult to construct an example which shows that the ordinal welfare factor of RSD is $O(1/K)$. On the other hand, by the time t agents arrive, at most Kt agents are disrupted, and so in a random permutation the $(t+1)$th agent is unhappy with probability $\leq \frac{(K+1)t}{n-t}$. Integrating, this gives that $\frac{n}{2K} - o(\frac{n}{K})$ agents are happy in expectation.

To conclude, in this paper we studied the social welfare of two well studied mechanisms, RSD and PS, for one-sided matching markets. We focussed on two measures: one was the ordinal welfare factor, and the other was the linear utilities measure. We performed a tight analysis of the ordinal welfare factors of both mechanisms, and the linear welfare factor in the case of efficient instances. An open problem is to perform a tighter analysis of linear welfare factor in general instances.We think the notion of ordinal welfare factor will be useful for other problems as well where the utilities are expressed as preference lists rather than precise numbers. Examples which come to mind are scheduling, voting, and ranking.

References

1. Abdulkadiroglu, A., Sönmez, T.: Ordinal efficiency and dominated sets of assignments. Journal of Economic Theory 112, 157–172 (2003)
2. Abdulkadiroglu, A., Sönmez, T.: Matching markets: Theory and practice. Prepared for the Econometric Society World Congress, China (2010), available at first author's website
3. Abraham, D., Irving, R., Kavitha, T., Melhorn, K.: Popular matchings. SIAM Journ. of Computing 37, 1030–1045 (2007)
4. Aggarwal, G., Goel, G., Karande, C., Mehta, A.: Online vertex-weighted bipartite matching and single-bid budgeted allocations. In: SODA (2011)

5. Ashlagi, I., Fischer, F., Kash, I., Procaccia, A.D.: Mix and Match. In: Proceedings of the ACM Conference on Electronic Commerce, EC 2010 (2010)
6. Bhalgat, A., Chakrabarty, D., Khanna, S.: Social welfare in one-sided matching markets without money. Technical report, http://arxiv.org/abs/1104.2964
7. Birnbaum, B., Mathieu, C.: On-line bipartite matching made simple. ACM SIGACT News 39 (2008)
8. Bogomolnaia, A., Moulin, H.: A new solution to the random assignment problem. Journal of Economic Theory 100, 295–328 (2001)
9. Bogomolnaia, A., Moulin, H.: A simple random assignment problem with a unique solution. Economic Theory 19, 623–635 (2002)
10. Budish, E., Che, Y.K., Kojima, F., Milgrom, P.: Designing random allocation mechanisms: Theory and applications. Unpublished Manuscript available at first author's personal webpage (October 2010)
11. Cres, H., Moulin, H.: Scheduling with opting out: Improving upon random priority. Operations Research 49, 565–577 (2001)
12. Dughmi, S., Ghosh, A.: Truthful assignment without money. In: Proceedings of the 11th ACM Conference on Electronic Commerce, EC 2010 (2010)
13. Gale, D., Shapley, L.: College admissions and the stability of marriage. American Mathematical Monthly 69, 9–15 (1962)
14. Gardenfors, P.: Match making: assignments based on bilateral preferences. Behavorial Sciences 20, 166–173 (1975)
15. Goel, G., Mehta, A.: Online budgeted matching in random input models with applications to adwords. In: Symposium on Discrete Algorithms, SODA (2008)
16. Guo, M., Conitzer, V.: Strategy-proof allocation of multiple items between two agents without payments or priors. In: 9th International Joint Conference on Autonomous Agents and Multi Agent Systems, AAMAS 2010 (2010)
17. Huang, C.-C., Kavitha, T., Michail, D., Nasre, M.: Bounded unpopularity matchings. In: Proceedings of the 11th Scandinavian Workshop on Algorithm Theory (2008)
18. Hylland, A., Zeckhauser, R.: The efficient allocation of individuals to positions. American Mathematical Monthly 69, 9–15 (1962)
19. Karp, R., Vazirani, U., Vazirani, V.: An optimal algorithm for online bipartite matching. In: 22nd Annual ACM Symposium on Theory of Computing (1990)
20. Katta, A.K., Sethurman, J.: A solution to the random assignment problem on the full preference domain. Journal of Economic Theory 131, 231–250 (2006)
21. Kavitha, T., Mestre, J., Nasre, M.: Popular mixed matchings. Theoretical Computer Science, Article in Press (2010)
22. McCutchen, R.M.: The least-unpopularity-factor and least-unpopularity-margin criteria for matching problems with one-sided preferences. In: Laber, E.S., Bornstein, C., Nogueira, L.T., Faria, L. (eds.) LATIN 2008. LNCS, vol. 4957, pp. 593–604. Springer, Heidelberg (2008)
23. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. In: Proceedings of the 10th ACM Conference on Electronic Commerce, EC 2009 (2009)
24. Roth, A.E., Sotomayor, M.: Two-Sided Matching: A Study on Game- Theoretic Modelling. Cambridge University Press, Cambridge (1990)
25. Sönmez, T., Ünver, U.: Matching, allocation, and exchange of discrete resources. In: Handbook of Social Economics, forthcoming, available at second author's website
26. Svensson, L.: Strategyproof allocation of indivisible goods. Social Choice and Welfare 16, 557–567 (1999)
27. Zhou, L.: On a conjecture by gale about one-sided matching problems. Journal of Economic Theory 52, 123–135 (1990)