

A Truthful Cardinal Mechanism for One-Sided Matching*

Rediet Abebe [†] Richard Cole [‡] Vasilis Gkatzelis [§] Jason D. Hartline [¶]

Abstract

We revisit the well-studied problem of designing mechanisms for one-sided matching markets, where a set of n agents needs to be matched to a set of n heterogeneous items. Each agent i has a value $v_{i,j}$ for each item j , and these values are private information that the agents may misreport if doing so leads to a preferred outcome. Ensuring that the agents have no incentive to misreport requires a careful design of the matching mechanism, and mechanisms proposed in the literature mitigate this issue by eliciting only the *ordinal* preferences of the agents, i.e., their ranking of the items from most to least preferred. However, the efficiency guarantees of these mechanisms are based only on weak measures that are oblivious to the underlying values. In this paper we achieve stronger performance guarantees by introducing a mechanism that truthfully elicits the full *cardinal* preferences of the agents, i.e., all of the $v_{i,j}$ values. We evaluate the performance of this mechanism using the much more demanding Nash bargaining solution as a benchmark, and we prove that our mechanism significantly outperforms all ordinal mechanisms (even non-truthful ones). To prove our approximation bounds, we also study the population monotonicity of the Nash bargaining solution in the context of matching markets, providing both upper and lower bounds which are of independent interest.

1 Introduction

In this paper we consider the classic “house allocation” problem of Hylland and Zeckhauser (1979). A set of n agents are to be matched, one-to-one, to a set of n items and each agent i has a value $v_{i,j}$ for each item j . A randomized matching mechanism outputs a probability distribution over matchings, which corresponds to a doubly-stochastic matrix p , providing the probability $p_{i,j}$ that i will be matched to j ; the expected utility of i in p is $\sum_j v_{i,j} p_{i,j}$. The goal in this setting is to generate a fair and efficient randomized matching, which crucially depends on the values of the agents, and the main obstacle is the fact that the $v_{i,j}$ values of each agent i are private information of this agent. Therefore, a successful mechanism needs to elicit the

agents’ preferences and output a desired matching; on the other hand, each agent’s goal is to maximize her expected utility, so an agent can strategically misreport her preferences if this increases her expected utility.

This tension between the objectives of the designer and those of the participants lies at the core of the sub-field of economics known as mechanism design. Most of the proposed solutions in the mechanism design literature, however, leverage monetary payments as the main tool that the designer can use to incentivize truthful reporting by the agents. This is in contrast to the matching mechanisms in the house allocation setting which cannot use such payments, making the mechanism design problem considered in this paper particularly demanding. In the absence of monetary payments, an alternative tool for simulating the impact of payments is to use “money burning” (Hartline and Roughgarden, 2008). In our setting, this could correspond to keeping some of the items unmatched with positive probability, thus penalizing the agents just like monetary payments would. But, in many settings, including the house allocation problem studied in this paper, this would be unacceptable (e.g., it would imply that some agents may remain homeless while some houses remain unoccupied).

Our main result is a novel application of random sampling that enables the use of known money burning mechanisms, while ensuring that every agent and item is matched. In other words, our technique takes advantage of the improved incentives that these money-burning mechanisms provide, but without suffering their most important drawback. To verify the usefulness of this approach, we combine it with the *partial allocation* (PA) mechanism of Cole et al. (2013), giving rise to a new mechanism that incentivizes the agents to always truthfully report their *cardinal* preferences, i.e., their $v_{i,j}$ values, and yields an outcome p that is approximately both fair and efficient.¹ We measure the performance of our mechanism using the canonical benchmark defined by the Nash bargaining solution and show that our mechanism outperforms the standard mechanisms with the same, or weaker, incentive properties.

*Rediet Abebe was supported in part by a Facebook scholarship. Richard Cole was supported in part by NSF grants CCF-1527568 and CCF-1909538. Vasilis Gkatzelis was supported in part by NSF grant CCF-1755955. Jason Hartline was supported by NSF grant CCF-1618502 and part of this work took place while he was visiting Harvard University.

[†]Society of Fellows, Harvard University

[‡]Courant Institute, New York University

[§]Computer Science Department, Drexel University

[¶]Computer Science Department, Northwestern University

¹Zhou (1990) shows that there is no mechanism that is truthful, symmetric, and Pareto efficient; thus, some notion of approximation is necessary.

The literature on one-sided matching has considered three main approaches, none of which gives rise to mechanisms that are both truthful and obtain a non-trivial approximation of the aforementioned benchmark. Hylland and Zeckhauser (1979) propose the *competitive equilibrium from equal incomes* (CEEI), which depends on the $v_{i,j}$ values in a non-trivial way, but it provides the agents with strong incentives to misreport these values, especially for small problem instances. The *random serial dictatorship* (RSD), or random priority, mechanism is an important mechanism with a long history in practice. This mechanism randomly orders the agents and, following this order, gives to each agent her favorite item among the ones that are still available. RSD is an *ordinal* mechanism: it requires only the ordinal preferences of each agent, i.e., only her ranking of the items from most to least preferred. It elicits this information truthfully, but its outcomes can be very inefficient. The *probabilistic serial* (PS) mechanism of Bogomolnaia and Moulin (2001) is another ordinal mechanism, and its outcome is computed by continuously allocating to each agent portions of her most preferred item that has not already been fully allocated. PS satisfies an ordinal notion of efficiency, but it achieves only a trivial approximation of our stronger benchmark, and it is not truthful. We provide a more detailed discussion regarding these mechanisms and other related work in Section 7.

Aiming to provide stronger efficiency and fairness guarantees compared to known mechanisms, we consider a cardinal benchmark: the well-studied *Nash bargaining solution*, proposed by Nash (1950). Given a *disagreement point*, i.e., the “status quo” that would arise if negotiations among the agents were to break down, the Nash bargaining solution is the outcome that maximizes the product of the agents’ marginal utilities relative to their utility for the disagreement point. This outcome indicates the utility that each agent “deserves”, so we use this utility as the benchmark for that agent. The choice of disagreement point can depend on the application at hand: if a buyer and a seller are negotiating a transaction, the disagreement point could be that the seller keeps the goods and the buyer keeps her money. In one-sided matching markets the disagreement point needs to be a matching because leaving an agent without a house is infeasible. Since all agents have symmetric claims on the items when entering the market, we let the disagreement point be a matching chosen uniformly at random, ensuring that each agent is equally likely to be matched to each item. The Nash bargaining solution therefore corresponds to the doubly-stochastic matrix p that maximizes $\prod_i (\sum_j v_{i,j} p_{i,j} - o_i)$, where $o_i = \frac{1}{n} \sum_j v_{i,j}$ is the expected utility of agent i for an item chosen uniformly at random.

Since no truthful and symmetric mechanism can guarantee Pareto efficiency (Zhou, 1990), it is clearly impossible for a truthful mechanism to implement the Nash bargaining solution, which is symmetric and Pareto efficient. Thus, we consider the problem of approximating this solution. Specifically, the Nash bargaining solution defines the utility that each agent deserves and our goal is to ensure that every agent receives a good approximation of that benchmark. Formally, a mechanism is a β -approximation if the utility of each agent is at least a β fraction of her utility in the Nash bargaining solution. Note that, once the valuations of each agent i are adjusted by subtracting o_i , then our objective corresponds to the *Nash social welfare* (NSW), which has recently received a lot of attention in the fair division literature (e.g., Cole and Gkatzelis, 2018; Garg et al., 2018; Caragiannis et al., 2016; Barman et al., 2018; Brânzei et al., 2017). The NSW maximizing outcome is *proportionally fair* in that it satisfies a multiplicative version of Pareto efficiency, namely, the utility of an agent cannot be increased by a multiplicative factor without decreasing the product of utilities of other agents by a greater multiplicative factor.

En route to proving our mechanism’s approximation bounds, we also provide an analysis of the Nash bargaining solution with respect to its *population monotonicity*, which is of independent interest. It has long been known that, unlike the Kalai-Smorodinsky solution, the Nash bargaining solution can violate population monotonicity for some instances of the bargaining problem (Thomson, 1983; Thomson and Lensberg, 1989). That is, there exist instances where removing some of the agents and computing the updated Nash bargaining solution can decrease the utility of some of the remaining agents. When allocating items among competing agents, this lack of monotonicity is somewhat counterintuitive. Why would the decreased competition from agents departing the market not lead to (weakly) increased utility for the agents remaining in the market? Indeed, we show that population monotonicity can be violated in the Nash bargaining solution for matching markets. Effectively, the constraint that the allocation is a distribution over perfect matching introduces positive externalities between agents.

In order to quantify the extent to which one of the remaining agents’ utility can drop after such a change in the agent population, the bargaining literature in economics introduced the *opportunity structure* notion (e.g., see the book by Thomson and Lensberg, 1989, and references therein). This structure identifies the largest factor by which a remaining agent’s utility can drop after some subset of agents is removed. In fact, resembling the standard computer science approach, the

opportunity structure is defined as the *worst-case* factor over all instances, all removed subsets of agents, and all remaining agents. In this paper we provide essentially tight upper and lower bounds for this factor in the context of matching markets, showing that in carefully designed worst-case instances, this factor can grow faster than a polylogarithmic function of the number of agents, yet slower than any polynomial. Apart from the broader interest in understanding this measure in matching markets, we show that the upper bound on the population non-monotonicity provides, up to constant factors, an upper bound on the approximation factor of the truthful matching mechanism that we define.

Our Results. In this paper we introduce a random sampling technique which allows us to translate non-trivial truthful one-sided matching mechanisms that may produce *partial* matchings (i.e., possibly leaving some agents unmatched) into ones where (i) every agent is always assigned an item, and (ii) the incentives for truthful reporting of preferences are maintained. For example, the truthfulness guarantee of the PA mechanism of Cole et al. (2013) depends heavily on its ability to penalize the agents that cause inconvenience to others; it thereby ensures that none of these agents is misreporting their preferences. Since monetary payments are prohibited, this mechanism penalizes the agents by assigning positive probability to outcomes that leave them unmatched. Such partial matchings, however, are unacceptable in the house allocation problem. Every agent, no matter what values she reports, needs to be guaranteed an item, and this constraint significantly restricts our ability to introduce penalties. Nevertheless, we show that we can still recreate such penalties by using random sampling. Applying our sampling technique to the PA mechanism, we define the *randomized partial improvement* (RPI) mechanism, which significantly outperforms all the standard matching mechanisms with respect to the Nash bargaining benchmark.

In essence, the RPI mechanism endows agents with a baseline allocation given by a uniformly random item and then uses the PA mechanism to improve the agents' utility relative to this baseline. In reality, it is not possible to simultaneously maintain the baseline and offer improvements to all agents, so RPI circumvents this impossibility by imposing these two conditions on a sample of half of the agents instead. With half the agents (but all of the items) there is sufficient flexibility to faithfully implement the PA mechanism with the outside option of a uniform random house. After finalizing the allocation of the sampled agents, RPI then recursively allocates the unallocated portions of the items to the remaining agents.

As an intermediate step toward the theoretical analysis of RPI's approximation factor, we study the extent to which population monotonicity may be violated in a one-sided matching market instance. We refer to an instance as ρ -utility monotonic if removing a subset of its agents can *decrease* a remaining agent's utility in the new Nash bargaining solution by a factor no more than ρ . We show that, for a very carefully constructed (and somewhat contrived) family of instances, ρ can be as high as $\Omega(2^{\sqrt{\log n}/2})$ and we complement this bound with an essentially tight upper bound, by proving that for any one-side matching instance ρ is no more than $O(2^{2\sqrt{\log n}}) \subseteq o(n^\epsilon)$ for any constant $\epsilon > 0$.

Apart from the broader interest in understanding the extent to which the Nash bargaining solution may violate population monotonicity, our upper bound on ρ also directly implies an upper bound for the approximation factor of RPI. Specifically, we prove that RPI guarantees to *every* agent a $4e\rho$ approximation of the utility that she gets in the Nash bargaining benchmark. Therefore, as a corollary, we conclude that RPI approximates the Nash bargaining benchmark within $O(2^{2\sqrt{\log n}}) \subseteq o(n^\epsilon)$ for any constant $\epsilon > 0$, even with the worst case choice of ρ . In stark contrast to this upper bound, which is strictly better than any polynomial, we show that the approximation factor of all ordinal mechanisms (even ones that are not truthful, such as probabilistic serial) grows linearly with the number of agents. Therefore, our mechanism significantly outperforms all ordinal mechanisms while at the same time satisfying truthfulness.

Structure. Section 2 provides some preliminary definitions and Section 3 formally introduces the benchmark and approximation measure used throughout the paper. Our results showing that ordinal mechanisms fail to achieve any non-trivial approximation are in Section 4; Section 5 includes the description of our mechanism and the proofs regarding its truthfulness and fairness guarantees. Finally, in Section 6 we study the population monotonicity of the Nash bargaining solution and provide both upper and lower bounds for it.

2 Preliminaries

Given a set N of n agents and a set M of n items, a randomized matching can be represented by a doubly-stochastic matrix p of marginal probabilities, where $p_{i,j}$ denotes the marginal probability that agent i is allocated item j . Clearly, any probability distribution over matchings implies a double-stochastic matrix, and the Birkhoff-von-Neumann theorem shows that *any* doubly-stochastic matrix can be implemented as a probability distribution over matchings. Denote by v a matrix of agent values where $v_{i,j}$ is the value of agent i for item

j . The expected utility of agent i for random matching p is $u_i = \sum_{j \in M} v_{i,j} p_{i,j}$. The random matching p that a mechanism outputs when the agents' reported values are v is denoted by $p(v)$.

For each agent i , her values $v_i = (v_{i,1}, \dots, v_{i,n})$ are private and a matching mechanism must be designed to properly elicit them. A mechanism is *truthful* if it is a dominant strategy for each agent i to report her true values. That is, if we let $p(w_i, v_{-i})$ denote the outcome of the mechanism when agent i reports values w_i and all the other agents report values v_{-i} , then a mechanism is truthful if for every agent i , any matrix of values v , and any misreports w_i :

$$\sum_{j \in M} v_{i,j} p_{i,j}(v) \geq \sum_{j \in M} v_{i,j} p_{i,j}(w_i, v_{-i}).$$

Our benchmark, formally defined in the following section, uses the *Nash social welfare* (NSW) objective on appropriately adjusted agent valuations. The NSW maximizing outcome is known to provide a balance between fairness and efficiency by maximizing the geometric mean (or, equivalently, the product) of the agents' expected utilities, i.e., $\max_p \prod_i \left(\sum_j v_{i,j} p_{i,j} \right)$. The partial allocation mechanism from Cole et al. (2013) provides a truthful approximation of that outcome and can be easily adapted to randomized matchings by interpreting fractional allocations as probabilities.

DEFINITION 1. *The partial allocation (PA) mechanism on values v works as follows:*

1. Compute the doubly-stochastic matrix $p^{\text{NSW}}(v)$ that maximizes the Nash social welfare.
2. For each agent i , compute f_i as follows:
 - (a) Let u_k be agent k 's utility in $p^{\text{NSW}}(v)$.
 - (b) Let u'_k be agent k 's utility in $p^{\text{NSW}}(v_{-i})$, i.e., in the NSW maximizing allocation with agent i absent and all other agents restricted to one unit, i.e., $\sum_j p_{k,j}^{\text{NSW}}(v_{-i}) = 1$ for all $k \neq i$.
 - (c) Let $f_i = \prod_{k \neq i} u_k / \prod_{k \neq i} u'_k$.
3. Allocate each item j to each agent i with probability $q_{i,j} = f_i p_{i,j}^{\text{NSW}}(v)$.

Note that the fraction f_i of the NSW maximizing assignment allocated to agent i is equal to the relative loss in utility that i 's presence imposes on the other agents. The denominator is independent of i 's declared valuations, so, in maximizing $f_i \cdot u_i$, which would be agent i 's goal, she is maximizing the NSW when she reports truthfully. Cole et al. (2013) show that $f_i \in (1/e, 1]$, without the unit constraint on allocations, but the same argument holds with the unit constraint.

THEOREM 2.1. (COLE ET AL., 2013) *The partial allocation mechanism is truthful, feasible, and allocates each agent i at fraction f_i of the NSW maximizing assignment, where f_i is at least $1/e$.*

3 The Nash Bargaining Benchmark

In this section, we define our cardinal benchmark as well as an approximation measure for evaluating mechanisms for the one-sided matching problem. Our benchmark is the Nash bargaining solution with a uniformly random matching as the disagreement point. Each agent i 's expected utility for this disagreement point is $o_i = \frac{1}{n} \sum_j v_{i,j}$ and the Nash bargaining solution is the outcome p^* that maximizes the Nash Social Welfare objective with respect to the marginal valuations $v - o$. In other words, the Nash bargaining solution distributes the additional value, beyond each agent's outside option, in a fair and efficient manner.

DEFINITION 2. *The Nash bargaining solution with disagreement point $(o_i)_{i \in N}$ is*

$$p^* = \operatorname{argmax}_p \left\{ \prod_i \left(\sum_j v_{i,j} p_{i,j} - o_i \right) \right\},$$

where every agent i is constrained to have non-negative utility $\sum_j v_{i,j} p_{i,j} - o_i \geq 0$.

Apart from its fairness properties, this benchmark is also appealing because of its invariance to additive shifts and multiplicative scalings of any agent's values for the items. Shifting all the values of an agent by adding some constant does not affect the marginal values after the outside option is subtracted. Also, scaling all of the values of an agent by some constant does not have any impact on what the Nash bargaining solution, p^* , is; the product value of every outcome is multiplied by the same constant, and hence the optimum is unaffected. As a result, we do not need to assume that the values reported by the agents are scaled in any particular way. One thing to note about the benchmark being invariant to these changes is that, on instances where the agents' values are identical up to shifts and scales, the benchmark assignment is the uniform random assignment.²

²The combined property of shift and scale invariance has some counterintuitive implications. Consider an example instance where all agents i have value $v_{i,1} > 1$ for item 1, and $v_{i,j} = 1$ for all other items $j \in \{2, \dots, n\}$. In the Nash bargaining solution, all agents receive a uniform random item and in particular a $1/n$ fraction of the preferred item 1. This outcome may seem surprising as it does not account for the possibility that some agents may prefer item 1 much more than other agents. This uniform outcome results because the agents' preferences are equivalent up to additive and multiplicative shifts.

Our goal is to approximate p^* , the Nash bargaining solution with disagreement point $(o_i)_{i \in N}$, using the following per-agent approximation measure.

DEFINITION 3. *The per-agent approximation of mechanism p with respect to benchmark assignment p^* is the worst-case ratio of the utility of any agent in p^* and p ,*

$$\max_v \left\{ \max_i \left\{ \frac{\sum_j v_{i,j} p_{i,j}^*(v)}{\sum_j v_{i,j} p_{i,j}(v)} \right\} \right\}.$$

4 Inapproximability by Ordinal Mechanisms

Ordinal mechanisms are popular in the literature on matching. Rather than asking agents for cardinal values for each item, an ordinal mechanism need only solicit an agent's preference order over the items. Two prevalent ordinal mechanisms are the random serial dictatorship (RSD) and probabilistic serial (PS) mechanisms. One of our main motivations for studying cardinal mechanisms in this paper is that ordinal mechanisms are bound to generate unfair allocations for some instances, due to the fact that they disregard the intensity of the agents' preferences; even when the agents agree, or partially agree, on their preference order, they may still disagree on preference intensities. A mechanism that does not take these intensities into consideration is, for example, unable to distinguish between agents whose favorite item is very strongly preferred over the rest, and agents who have only a slight preference for their top item over the rest.

Our first lower bound shows that the random serial dictatorship mechanism can be very unfair to some agent, leading to an approximation factor as bad as n (the number of agents).

LEMMA 4.1. *The worst case approximation ratio of the random serial dictatorship (RSD) mechanism with respect to the Nash bargaining benchmark is n .*

Proof. Consider the example where agent 1 has value 1 for item 1 and no value for any other item, and each agent $i \geq 2$ has value 1 for item 1, value $1 - \epsilon$ for item i , and no value for other items:

$$v = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 \\ 1 & 1 - \epsilon & 0 & \dots & 0 \\ 1 & 0 & 1 - \epsilon & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & \dots & 1 - \epsilon \end{bmatrix}.$$

In RSD an ordering of the agents is generated uniformly at random, and then each agent is allocated her favorite available item in that order. In this instance, the first agent in the random ordering will always select item 1,

and every agent has the same probability, $1/n$, of being ordered first. Since agent 1 has no value for any other item, the expected utility of this agent 1 in RSD is $1/n$.

On the other hand, as ϵ approaches zero, the Nash bargaining solution assigns each agent i to item i with probability that approaches 1. To verify this fact, note that for $\epsilon = 0$ the Nash bargaining would assign agent i to item i with probability 1, and observe that the distribution that RSD outputs is continuous in ϵ . Thus, the utility of each agent in the Nash bargaining solution – and specifically of agent 1 – approaches 1. As a result, the RSD mechanism is being unfair to agent 1, leading to an approximation factor of n . \square

In fact, with a small modification of the instance used to verify how unfair the RSD mechanism can be, the following theorem shows that *every* ordinal mechanism is susceptible to this issue.

THEOREM 4.1. *The worst case approximation ratio of any ordinal mechanism to the Nash bargaining benchmark is at least $n - 1$.*

Proof. Consider the following instance v , where agents correspond to rows and items to columns:

$$v = \begin{bmatrix} 1 & \epsilon & 0 & 0 & \dots & 0 & 0 \\ 1 & 0 & 1 - \epsilon & 0 & \dots & 0 & 0 \\ 1 & 0 & 0 & 1 - \epsilon & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & 0 & 0 & 0 & \dots & 1 - \epsilon & 0 \\ 1 & 0 & 0 & 0 & \dots & 0 & 1 - \epsilon \\ 0 & 1 & 1 & 1 & \dots & 1 & 1 \end{bmatrix}.$$

A key property of this instance is that the top $n - 1$ agents are ordinally indistinguishable. Each of them ranks item 1 first, one of items $\{2, \dots, n\}$ second, and all other items last. On the other hand, each item $j \in \{2, \dots, n\}$ is ordinally indistinguishable. Each is ranked second by exactly one of the top $n - 1$ agents and ranked equivalently by agent n .

Fix an ordinal mechanism. The ordinal indistinguishability of agents $\{1, \dots, n - 1\}$ implies, without loss of generality up to agent relabeling, that agent 1 receives item 1 with probability at most $1/(n - 1)$. Thus, in the limit of ϵ going to 0, agent 1 obtains a utility of $1/(n - 1)$ in this ordinal mechanism.

The Nash bargaining solution is continuous in ϵ and with $\epsilon = 0$ it gives each agent the maximum utility of 1 by allocating item 1 to agent 1, item 2 to agent n , and item $i + 1$ to agent i for $i \in \{2, \dots, n - 1\}$. Thus, in the limit, as ϵ goes to zero the Nash bargaining solution gives agent 1 a utility of 1. Therefore, the per-agent approximation of any ordinal mechanism with respect to the Nash bargaining benchmark is $n - 1$. \square

5 Randomized Partial Improvement

In this section, we define the random partial improvement matching mechanism. This mechanism truthfully elicits the agents' cardinal preferences and uses them in a non-trivial manner to select an outcome. We prove that the per-agent approximation of this mechanism with respect to the Nash bargaining benchmark is proportional to the population monotonicity of the benchmark and its worst-case approximation is significantly better than that of any ordinal mechanism. The approach of the mechanism is to run the PA mechanism with the outside option given by the uniform random assignment on a large sample of the agents and a large fraction of the supply. The resulting mechanism inherits the truthfulness of the PA mechanism.

There are two key difficulties with this approach. First, in order to faithfully implement the outside option, some of the supply needs to be kept aside in the same proportion as the original supply. To enable this set aside, we need to reduce the allocation consumed by the PA mechanism; we achieve this with a novel use of random sampling (cf. Goldberg et al., 2006). Second, it is non-trivial to compare an agent's utility across Nash social welfare maximizing assignments for the original market and a sample of the market. A major endeavor of our analysis shows that per-agent utility is approximately monotone, i.e., the fraction of an agent's utility that is lost as the competition from other agents decreases is non-trivially bounded. (Note, competition from other agents decreases as they are removed from the market.) Our mechanism, then, is structured to take advantage of this approximate monotonicity.

The mechanism is defined by a sequence of steps that gradually construct a doubly-stochastic matrix. By the Birkhoff von Neumann theorem, this matrix corresponds to a probability distribution over matchings. The high-level steps and intuition are as follows: the mechanism samples half the agents and runs at half scale (i.e., with half-unit-demand agents and half-unit-supply items) the PA mechanism with the outside option given by the uniform random assignment. The total demand of half the agents (roughly $n/2$) with half-unit demand is a quarter of the total supply (roughly $n/4$), so there is a leftover $n/4$ of supply from the half-units on which the PA mechanism was run. A further one-quarter of each of the n units is used to provide a half-unit of the outside option to each of the (roughly) $n/2$ agents in the sample. The final quarter is used to replace as necessary the fractions of items withheld due to the fractional reduction in the PA mechanism. Necessarily, the one-unit allocation to these agents uses up half the supply. The remaining half of the supply is then allocated recursively to the remaining half of the agents.

A formal description of this mechanism is below. Since we call this mechanism recursively after some agents' allocation has been finalized (in the form of marginal probabilities) and some portions of the items have been allocated, we define it for the remaining $\bar{n} \leq n$ agents and the original n items whose capacities may have been reduced from 1 to $(c_i)_{i \in M} \in [0, 1]^n$.

DEFINITION 4. *Given some value $n_0 \in \mathbb{N}$, the randomized partial improvement (RPI) mechanism on $\bar{n} \leq n$ agents and n items with supplies c_1, \dots, c_n such that $\sum_{j=1}^n c_j = \bar{n}$ works as follows:*

1. *If $\bar{n} < n_0$, allocate the remaining item capacities uniformly at random, i.e., return $p_{i,j} = c_j/\bar{n}$ for each agent i and terminate. Otherwise, continue.*
2. *Randomly sample a subset N' of $n' = \lceil \bar{n}/2 \rceil$ agents.*
3. *On the sampled agents run the PA mechanism with the outside option given by the uniform random assignment $o'_i = \frac{1}{\bar{n}} \sum_{j=1}^n v_{i,j} c_j$ from the supplies. Denote the allocation of item j to agent i by $q'_{i,j}$; the total amount allocated to agent i is $f'_i = \sum_{j=1}^n q'_{i,j}$.*
4. *Allocate to each $i \in N'$ half of their PA assignment and "pad" it with the outside option to ensure a unit allocation. As a result, the total allocation of item j to agent i is $p'_{i,j} = q'_{i,j}/2 + (1 - f'_i/2) c_j/\bar{n}$.*
5. *Recursively run RPI on the remaining $n'' = \bar{n} - n'$ agents and item supplies $c''_j = c_j - \sum_{i \in N'} p'_{i,j}$.*
6. *Return the assignment p that combines the assignment p' for N' and the assignment p'' returned by the recursive call for the remaining n'' agents.*

The following proof of correctness (feasibility and truthfulness) formalizes the intuition preceding the definition of the mechanism. The ideas of the proof are more transparent in the case where \bar{n} is even and, in particular, $\lceil \bar{n}/2 \rceil = \lfloor \bar{n}/2 \rfloor = \bar{n}/2$.

THEOREM 5.1. *The randomized partial improvement mechanism with $n_0 \geq 4$ on n unit-demand agents and n unit-supply items is feasible, i.e., it gives fractional allocations that produce a doubly stochastic matrix, and truthful, i.e., it is a dominant strategy for each agent to truthfully report her value for each item.*

Proof. Feasibility is proved by induction on the recursive definition of the mechanism. The inductive hypothesis is that the fractional allocation on \bar{n} agents with supplies c_1, \dots, c_n such that $\sum_j c_j = \bar{n}$ has a total fractional allocation to each agent of one, i.e., $\sum_j p_{i,j} = 1$ for each i , and a total fractional allocation of each item

equal to its supply, i.e., $\sum_i p_{i,j} = c_j$ for each j . The base case of $\bar{n} < n_0$ clearly satisfies the inductive hypothesis. For the inductive step, the key point to argue is that the supply c_j of each item j is sufficient to cover the allocation to the sampled agents.

This can be seen as follows. The $\lceil \bar{n}/2 \rceil$ sampled agents are allocated half of their PA assignment on the supplied capacities. Since $f'_i \geq \frac{1}{e}$ for all i by Theorem 2.1, this means that the amount of each agent's half PA assignment is $\sum_j q'_{i,j}/2 = f'_i/2 \geq 1/(2e)$. To ensure that each agent gets exactly one item in expectation, Step 4 pads this allocation with a uniformly random assignment, which may thus require up to $1 - 1/(2e)$ units for each of the $\lceil \bar{n}/2 \rceil$ sampled agents, for a total of $(1 - 1/(2e))\lceil \bar{n}/2 \rceil$. But, since the full PA assignment allocated no more than c_j of each item j , the half PA assignment set aside at least $c_j/2$ of each item, leading to a total of $\sum_j c_j/2 = \bar{n}/2$. We conclude by observing that $(1 - 1/(2e))\lceil \bar{n}/2 \rceil$ is at most $\bar{n}/2$ when $\bar{n} \geq n_0 = 4$, so the amount set aside from each item is sufficient to cover the sampled agents' allocation.

Truthfulness follows by considering each agent conditioned on the state of the mechanism during the recursive step where that agent is selected in the sampled set N' . The agent's report plays no role in determining the state at this point. Given the state, the outcome for this agent is fully determined by the PA mechanism which is truthful. Thus, the mechanism is truthful. \square

To bound the per-agent utility of the RPI mechanism, we analyze the contribution to the utility of an agent who is sampled in the outermost recursive call of the mechanism. An agent is sampled as such with probability at least one half, and otherwise the agent's utility is at least zero. The utility of these sampled agents is easily compared to the utility of the PA mechanism (without the agents that are not sampled). An issue significantly complicating the analysis of the approximation is the fact that we need to compare the utility of an agent sampled in this invocation of the PA mechanism with their utility in the Nash bargaining solution on the full set of agents. Counterintuitively, it is not true that these agents are always better off without the competition from the agents that are not sampled: there are instances where removing some of the competition, in fact, lowers the utility of an agent.

In Section 6 we define the ρ -utility monotonicity for NSW to be the maximum non-monotonicity of utility of any agent and sets of agents N and subset N' with NSW maximizing solutions p and p' respectively:

$$\rho := \max_{N' \subseteq N} \left\{ \max_{i \in N'} \left\{ \frac{\sum_j v_{i,j} p_{i,j}}{\sum_j v_{i,j} p'_{i,j}} \right\} \right\}.$$

This parameter quantifies the extent to which some agent may be worse off in the NSW maximizing solution after the removal of some subset of agents. Defining the worst-case value of ρ across instances and subsets as ρ^* , Section 6 bounds ρ^* between $\Omega(2^{\sqrt{\log n}/2})$ and $O(2^{2\sqrt{\log n}})$, the latter of which is $o(n^\epsilon)$ for any constant $\epsilon > 0$. It is worth noting that we ran experiments on a large set of instances and found that this value was actually no more than 1 in all of these instances.

THEOREM 5.2. *The randomized partial improvement mechanism with $n_0 = 4$ on n unit-demand agents and n unit-supply items is a $4e\rho$ approximation to the Nash bargaining solution with disagreement point given by the uniform random assignment.*

Proof. If $n < n_0 = 4$ then the base case of RPI is invoked and a uniform assignment is returned. With $n < 4$, however, this assignment is a $3 < 4e\rho$ approximation, as each agent obtains $1/3$ of each item.

Otherwise, we analyze the contribution to the utility of an agent conditioned on the agent being sampled in the first recursive call of the algorithm. This event happens with probability at least $1/2$. When this happens the utility of the agent is half the utility of PA on the sampled agents plus half the utility from the outside option. The ρ -utility monotonicity property implies that the utility of an agent in the NSW maximizing outcome on the sample is a ρ approximation to the same agent's utility in the NSW maximizing outcome on the full set of agents. Running PA guarantees an e fraction of this utility. Combining these steps and the fact that agents who are not sampled in the first recursive call still receive nonnegative utility, we obtain a $4e\rho$ approximation. \square

Combined with Theorem 6.1 from Section 6 which bounds ρ^* by $O(2^{2\sqrt{\log n}})$, we have the following.

COROLLARY 5. *The randomized partial improvement mechanism with $n_0 = 4$ on n unit-demand agents and n unit-supply items guarantees an approximation of the Nash bargaining solution with uniform outside option with a factor $O(2^{2\sqrt{\log n}}) \subseteq o(n^\epsilon)$ for any constant $\epsilon > 0$.*

6 Approximate Utility Monotonicity

A factor significantly complicating the analysis of the approximation of the random partial improvement mechanism is the fact that the benchmark is computed based on the Nash social welfare maximizing solution when all agents in N are present, while the mechanism's performance depends on the solution for N' , the sampled agents. The NSW maximizing solution for N and N' can generally be quite different. Moreover, as

it turns out, there are instances where the utilities of some agents in the NSW maximizing solution are non-monotone with respect to removal of other agents, i.e., there exist instances that exhibit positive externalities between agents. Table 1 gives a simple example of such an instance (discussed in detail later on) and the remainder of the section develops upper and lower bounds on the worst-case non-monotonicity of utility.

DEFINITION 6. *A matching environment on agents N is ρ -utility monotone if for any subset N' of N and any $i \in N'$ the utility of i in the NSW maximizing assignment, p' , for N' is at least a ρ approximation to the NSW maximizing assignment, p , for N :*

$$\rho := \max_{N' \subseteq N} \left\{ \max_{i \in N'} \left\{ \frac{\sum_j v_{i,j} p_{i,j}}{\sum_j v_{i,j} p'_{i,j}} \right\} \right\}.$$

This parameter quantifies the extent to which some agent may be worse off in the NSW solution after the removal of some subset of agents. We let ρ^* , denote the worst case value of ρ across instances; this value is known as the *opportunity structure* of the Nash bargaining solution for this class of instances (Thomson and Lensberg, 1989). In Section 6.1 we prove an upper bound of $O(2^{2\sqrt{\log n}})$, which is $o(n^\epsilon)$ for any constant $\epsilon > 0$, for the value of ρ^* over all one-sided matching instances, and in Section 6.2 we complement this result by proving a lower bound of $\Omega(2^{\sqrt{\log n}/2})$ for this value.

6.1 Upper Bound. Given a valuation matrix v and a random matching p , we henceforth use $u_i(p)$ to denote the expected utility of agent i for p given v , i.e., $\sum_{j \in M} v_{i,j} p_{i,j}$ (similarly, we use $u'_i(p)$ for valuation matrix v'). Note that, as we discussed in Section 3, the Nash bargaining solution is scale invariant. Therefore, if we scale the valuations of each agent i by some constant $c_i > 0$, then the Nash bargaining solution with respect to valuations $c_i v_{i,j}$ instead of $v_{i,j}$ will remain the same. This means that given some problem instance that yields a doubly stochastic matrix p as its Nash bargaining solution, we can always “normalize” the valuations of the agents so that every agent’s expected utility for p is equal to 1, and p remains that Nash bargaining solution of the normalized instance. This is a convenient normalization that we make use of below.

In order to prove the upper bound on ρ^* , we first prove the following very useful lemmata.

LEMMA 6.1. *Let p be a NSW maximizing solution, and v be the valuations normalized so that for every agent i , $u_i(p) = 1$. Then, if some agent i is allocated an item j with positive probability, i.e., $p_{i,j} > 0$, every other agent $k \neq i$ must have $v_{k,j} \leq v_{i,j} + 1$. Equivalently, $v_{i,j} \geq \max_{k \in N} \{v_{k,j}\} - 1$.*

Proof. For contradiction, assume that there exist two agents k and i and an item j such that $p_{i,j} > 0$ and $v_{k,j} = v_{i,j} + 1 + \delta$ for some $\delta > 0$. Since the expected utility of agent k is 1, there must also exist some item ℓ with $p_{k,\ell} > 0$ and $v_{k,\ell} \leq 1$ (otherwise the expected utility of agent k would be greater than 1). Note that $\ell \neq j$, since $v_{k,j} = v_{i,j} + 1 + \delta > 1$, whereas $v_{k,\ell} \leq 1$.

Let p' be a probability distribution identical to p , except $p'_{k,j} = p_{k,j} + \epsilon$, $p'_{i,j} = p_{i,j} - \epsilon$, $p'_{k,\ell} = p_{k,\ell} - \epsilon$, and $p'_{i,\ell} = p_{i,\ell} + \epsilon$, for some positive $\epsilon < \min\{p_{i,j}, p_{k,\ell}\}$, whose exact value we choose later on. In other words, p' swaps probability ϵ between agents i, k and items j, ℓ . The expected utility of agent k in p' is

$$\begin{aligned} u_k(p') &= 1 + \epsilon v_{k,j} - \epsilon v_{k,\ell} \\ &\geq 1 + \epsilon(v_{i,j} + 1 + \delta) - \epsilon \\ &= 1 + \epsilon v_{i,j} + \epsilon \delta, \end{aligned}$$

and the expected utility of agent i in p' is

$$\begin{aligned} u_i(p') &= 1 + \epsilon v_{i,\ell} - \epsilon v_{i,j} \\ &\geq 1 - \epsilon v_{i,j}. \end{aligned}$$

Since every other agent’s expected utility is the same in p' and p (equal to 1), the NSW of p' is

$$\begin{aligned} \prod_{i \in N} u_i(p') &\geq (1 + \epsilon v_{i,j} + \epsilon \delta)(1 - \epsilon v_{i,j}) \\ &= 1 + \epsilon(\delta - \epsilon(v_{i,j}^2 + \delta v_{i,j})). \end{aligned}$$

Therefore, if we let $\epsilon < \delta/(v_{i,j}^2 + \delta v_{i,j})$, the NSW of p' is greater than 1, which is the NSW of p , contradicting the fact that p is a NSW maximizing solution. \square

LEMMA 6.2. *Given a problem instance, let p and \bar{p} be the NSW maximizing outcomes before and after (respectively) some subset of the agents has been removed. If among the remaining agents there exists a set of agents N_1 and a constant $d \geq 12$ such that every agent $i \in N_1$ has $u_i(\bar{p}) \leq u_i(p)/d$, then there also exists a larger set N_2 of remaining agents such that $|N_2| \geq d|N_1|/3$ and for all agents $k \in N_2$ we have $u_k(\bar{p}) \leq 4u_k(p)/d$.*

Proof. Without loss of generality, let v be the agent valuations normalized so that $u_i(p) = 1$ for every agent i , and v' be the valuations normalized so that $u'_i(\bar{p}) = 1$. Given the $v_{i,j}$ values that yield $u_i(p) = 1$, we can get the $v'_{i,j}$ values that yield $u'_i(\bar{p}) = 1$ using the simple formula $v'_{i,j} = v_{i,j} \cdot \frac{u_i(p)}{u_i(\bar{p})}$. In other words, for each agent i who is worse-off in \bar{p} compared to p , i.e., $u_i(\bar{p}) < u_i(p)$, we scale all of that agent’s item values up by the same factor, $u_i(p)/u_i(\bar{p})$. In particular, for each agent $i \in N_1$ this means that $v'_{i,j} \geq d v_{i,j}$ for every item j .

	A	B	C	A	B	C	A	B	C
<i>a</i>	1	2	0	1	0	0	1/2	1/2	0
<i>b</i>	0	2	1	0	1	0	0	1/2	1/2
<i>c</i>	0	0	1	0	0	1			

i. Agent valuations
ii. Initial solution
iii. Final solution

Table 1: A simple instance involving three agents; a, b, c ; and three items; A, B, C . The value of agent b in the Nash bargaining solution after the removal of agent c drops by a factor of $4/3$. The optimality of solutions (ii) and (iii) for Nash social welfare are intuitive; however, formal justification is given in Section 6.2 where this example is revisited.

For every $i \in N_1$ we know that the drop in that agent's value with respect to the original valuations is $u_i(p) - u_i(\bar{p}) \geq u_i(p)(1 - 1/d) = 1 - 1/d$. In order to account for that drop, we partition the set of items of which i is allocated more in p compared to \bar{p} into two sets depending on whether $v_{i,j} \geq 0.5$ or not: $M_i^h = \{j \in M : p_{i,j} > \bar{p}_{i,j} \text{ and } v_{i,j} \geq 0.5\}$ and $M_i^\ell = \{j \in M : p_{i,j} > \bar{p}_{i,j} \text{ and } v_{i,j} < 0.5\}$. We first show that from the aforementioned $1 - 1/d$ drop in value, no more than 0.5 could be due to the items in M_i^ℓ , since

$$\sum_{j \in M_i^\ell} (p_{i,j} - \bar{p}_{i,j})v_{i,j} < 0.5 \sum_{j \in M_i^\ell} (p_{i,j} - \bar{p}_{i,j}) \leq 0.5.$$

Therefore, at least $0.5 - 1/d$ of this drop in value for each agent $i \in N_1$ is due to items in M_i^h . Summing this up over all the agents in N_1 , we get

$$(6.1) \quad \sum_{i \in N_1} \sum_{j \in M_i^h} (p_{i,j} - \bar{p}_{i,j})v_{i,j} \geq \sum_{i \in N_1} \left(0.5 - \frac{1}{d}\right)$$

$$(6.2) \quad = |N_1| \left(0.5 - \frac{1}{d}\right).$$

Let $N_2 = \{k \in N : \bar{p}_{k,j} > 0 \text{ for some } j \in \bigcup_{i \in N_1} M_i^h\}$ be the set of agents that are allocated with positive probability in \bar{p} an item from M_i^h for some $i \in N_1$. Using Lemma 6.1 we get that for every item $j \in M_i^h$, if $\bar{p}_{k,j} > 0$ then $v_{k,j} \leq v_{i,j} + 1$ and $v'_{k,j} \geq v'_{i,j} - 1$. Using the fact that $v'_{i,j} \geq dv_{i,j}$ for every $i \in N_1$, shown above, the latter inequality also implies that $v'_{k,j} \geq dv_{i,j} - 1$. Therefore

$$\frac{v'_{k,j}}{v_{k,j}} \geq \frac{dv_{i,j} - 1}{v_{i,j} + 1} \geq d - \frac{d+1}{1.5} \geq \frac{d}{4},$$

where the second inequality uses the fact that $v_{i,j} \geq 0.5$ and the last inequality uses the fact that $d \geq 8$. This

implies that for every $k \in N_2$

$$\begin{aligned} u_k(\bar{p}) &= \sum_{j \in M} \bar{p}_{k,j} v_{k,j} \\ &\leq \sum_{j \in M} \bar{p}_{k,j} \frac{4}{d} v'_{k,j} \\ &= \frac{4}{d} \sum_{j \in M} \bar{p}_{k,j} v'_{k,j} \\ &= \frac{4u'_k(p)}{d} \\ &= \frac{4u_k(p)}{d}, \end{aligned}$$

where the last equation uses the fact that $u'_k(p) = u_k(p) = 1$ according to our normalization.

Since we have shown that for all $k \in N_2$ we have $u_k(\bar{p}) \leq 4u_k(p)/d$, it now suffices to show that the size of N_2 is at least $d|N_1|/3$. Since, for any item $j \in \bigcup_{i \in N_1} M_i^h$, any agent k with $\bar{p}_{i,j} > 0$ satisfies $v'_{k,j} \geq dv_{i,j} - 1$, the total value, with respect to valuations v' , generated by the item fractions of the items "lost" by the agents in N_1 is at least

$$\begin{aligned} &\sum_{i \in N_1} \sum_{j \in M_i^h} (p_{i,j} - \bar{p}_{i,j})(dv_{i,j} - 1) \\ &\geq d|N_1| \left(0.5 - \frac{1}{d}\right) - \sum_{i \in N_1} \sum_{j \in M_i^h} (p_{i,j} - \bar{p}_{i,j}) \\ &\geq d|N_1| \left(0.5 - \frac{1}{d}\right) - |N_1| \\ &\geq \frac{d-4}{2}|N_1| \\ &\geq \frac{d}{3}|N_1|, \end{aligned}$$

where the last inequality uses the fact that $d \geq 12$. But, since the total value of each agent in N_2 with respect to valuations v' is exactly 1, there need to be at least $\frac{d}{3}|N_1|$ agents in N_2 sharing this value, otherwise there would exist some agent $i \in N_2$ such that $u'_i(\bar{p}) > 1$. \square

THEOREM 6.1. *For any problem instance, the value of ρ is $O(2^{2\sqrt{\log n}}) \subseteq o(n^\epsilon)$ for any constant $\epsilon > 0$.*

Proof. In order to prove this bound, we will repeatedly apply the result of Lemma 6.2. Let p and \bar{p} be the NSW maximizing outcomes in a problem instance before and after some subset of the agents has been removed and, without loss of generality, let v be the agent valuations normalized so that $u_i(p) = 1$ for every agent i , and v' be the valuations normalized so that $u'_i(\bar{p}) = 1$.

By Definition 6, in an instance with utility monotonicity equal to ρ , there exists at least one agent $i \in N_1$ such that $u_i(p)/u_i(\bar{p}) = \rho$ or $u_i(\bar{p}) = u_i(p)/\rho$. If $\rho > 12$, then Lemma 6.2 would imply that there also exists a set N_2 of at least $\rho/3$ agents such that $u_k(\bar{p}) \leq 4u_k(p)/\rho = 4/\rho$ for every $k \in N_2$. Lemma 6.2, combined with the existence of the set N_2 , in turn, implies the existence of an even larger group N_3 of at least $(\frac{1}{3}\rho) \cdot (\frac{1}{3}\rho/4)$ agents, and each agent $k \in N_3$ has value $u_k(\bar{p}) \leq 16/\rho$. Applying Lemma 6.2 a total of α times thus implies the existence of a set of at least $(\rho/3)^\alpha \cdot (1/4)^{\alpha(\alpha-1)/2}$ agents such that each such agent k has value $u_k(\bar{p}) \leq 4^\alpha/\rho$. Assume that there exists some instance for which ρ is at least $4^{\sqrt{\log n}+1}$. If we choose $\alpha = \sqrt{\log n}$, however, this implies the existence of $(\rho/3)^\alpha \cdot (1/4)^{\alpha(\alpha-1)/2} \geq 4^{(\sqrt{\log n}+1) \cdot \sqrt{\log n}} / [(3/2)^{\sqrt{\log n}} \cdot 2^{\sqrt{\log n} \cdot \sqrt{\log n}}] \geq (8/3)^{\sqrt{\log n}} \cdot n \geq n$ agents of value at most $4^\alpha/\rho \leq 1/4$. But, this would imply that all the agents have a value less than 1 in \bar{p} , which contradicts the fact that \bar{p} is a NSW maximizing solution because the product in p is equal to 1. \square

6.2 Lower Bound. We conclude with a lower bound showing that for a very carefully designed (and somewhat artificial) family of instances, the upper bound of Theorem 6.1 is essentially tight.

THEOREM 6.2. *There exists a family of problem instances for which $\rho^* = \Omega(2^{\sqrt{\log n}/2})$.*

Due to space limitations and the complexity of the construction that yields Theorem 6.2, we defer its description to Appendix A. To exhibit how we use KKT conditions to prove that this elaborate construction implies the desired bound, we use the rest of this section to apply this approach to the much simpler construction of the example in Table 1, which yields a bound of $\rho^* \geq 4/3$.

Our lower bound construction in the appendix proceeds by building a family of instances (parameterized by the number of agents n), and in each instance, we define an “initial” setting in which all agents are present, and a “final” setting, in which some agents have been removed. For each setting, we identify the Nash bargaining solution, respectively called the initial and final

solution. We focus on a particular agent, called the *loser*, who is present in both settings. We show that the loser’s valuation drops by a multiplicative factor μ in going from the initial to the final solution, and consequently, $\rho \geq 1/\mu$ for that market and $\rho^* \geq 1/\mu$ overall.

To prove a lower bound, we need to be able to verify that a given doubly stochastic matrix is indeed the Nash bargaining solution of the instance at hand. We do so using the KKT conditions, which allow us to interpret these solutions as a form of market equilibrium. The optimization problem which yields the Nash bargaining solution in one-sided matching markets, is shown below (where $m = n$ is used to denote the number of items):

$$\begin{aligned} & \max \sum_{i=1}^n \log \left[\sum_{j=1}^n v_{i,j} p_{i,j} \right] \\ & \text{such that:} \quad \text{for all } i : \sum_{j=1}^m p_{i,j} \leq 1 \\ & \quad \quad \quad \text{for all } j : \sum_{i=1}^n p_{i,j} \leq 1 \\ & \quad \quad \quad \text{for all } i, j : p_{i,j} \geq 0. \end{aligned}$$

If t_j is the dual variable related to each item j , and q_i is the dual variable related to each agent i in the above program, then the KKT conditions state that:

$$(6.3) \quad \text{for all } j : t_j \geq 0, \quad \text{and } t_j > 0 \implies \sum_{i=1}^n p_{i,j} = 1$$

$$(6.4) \quad \text{for all } i : q_i \geq 0 \quad \text{and } q_i > 0 \implies \sum_{j=1}^m p_{i,j} = 1$$

$$(6.5) \quad \begin{aligned} & \text{for all } i, j : \frac{v_{i,j}}{t_j + q_i} \leq \sum_{j=1}^m v_{i,j} p_{i,j} \quad \text{and} \\ & p_{i,j} > 0 \implies \frac{v_{i,j}}{t_j + q_i} = \sum_{j=1}^m v_{i,j} p_{i,j} \end{aligned}$$

The KKT conditions are necessary and sufficient for the optimal solution when the constraints are linear and the objective is convex, as is the case here. To check whether a given candidate solution p is a Nash bargaining solution for some instance, we first normalize the valuations so that $\sum_{j=1}^m v_{i,j} p_{i,j} = 1$ for all i . Then, at a solution satisfying the KKT conditions we have $v_{i,j} = t_j + q_i$ if $p_{i,j} > 0$ and $v_{i,j} \leq t_j + q_i$ if $p_{i,j} = 0$. Thus a solution that satisfies these two conditions plus conditions (6.3)–(6.4) is a Nash bargaining solution. Based on this conditions, the values of t_j can be interpreted as item-specific “prices” and the values of q_i as agent-specific “prices”, leading to an interpretation

of the Nash bargaining solution as a market equilibrium: to “buy” a $p_{i,j}$ fraction of item j , agent i needs to spend $(t_j + q_i)p_{i,j}$, and each agent prefers to buy only items with the best value over price ratio (see condition (6.5)).

To illustrate the usefulness of these variables, which are used extensively in the appendix, we revisit the instance of Table 1 where the items are named A , B , and C ; the bidders a , b , and c ; and the unscaled valuations of the agents appear in Table 2(i).

First, we observe that in the initial equilibrium (with all agents present), a, b, c receive items A, B, C , respectively, each with probability 1. In Table 2(ii) we show the normalized values of the agents in this equilibrium and we also provide the dual variables t_j for each item j and q_i for each agent i . It is easy to verify that the aforementioned KKT conditions are satisfied in this case and hence this is indeed the Nash bargaining solution when all agents are present. If agent c is removed, then the final equilibrium finds a receiving each of A and B with probability $\frac{1}{2}$, while b receiving each of B and C with probability $\frac{1}{2}$. Table 2(iii) provides the scaled valuations and dual variable values for this outcome, and it is again easy to verify that KKT conditions are satisfied. In this example, bidder b is the loser. Using the valuations from Table 2(ii), her value in the initial equilibrium was 1 and it dropped to 0.75 in the final equilibrium, leading to $\rho = \frac{4}{3}$ in this example.

7 Further Related Work

Hylland and Zeckhauser (1979) study the problem of matching with cardinal preferences and the solution of competitive equilibrium from equal incomes (CEEI). CEEI gives both a natural cardinal notion of efficiency and of fairness. Recently, Alaei et al. (2017) give a polynomial time algorithm for computing the CEEI in matching markets when there are a constant number of distinct agent preferences. To our knowledge, the complexity of computing CEEI in general matching problems is unknown. With linear preferences, but without the unit-demand constraint, CEEI and Nash social welfare coincide and can be computed in polynomial time. Devanur and Kannan (2008) generalize this computational result to piecewise linear concave utilities when the number of goods is constant.

Recently, Budish (2011) considers the generalization from matching to a combinatorial assignment problem where agents may have non-linear preferences over bundles of goods, and shows that an approximate version of CEEI exists. This work also shows that, in large markets, the mechanism that outputs this approximate CEEI is asymptotically truthful. Heuristics for computing the CEEI outcome are given by Othman et al. (2010) and these heuristics have been deployed for the

course assignment problem by Budish et al. (2016). On the other hand, Othman et al. (2016) show that the computation of CEEI in these combinatorial assignment problems is PPAD-hard.

The Nash social welfare objective of our work compares to competitive equilibrium from equal incomes of the aforementioned works as follows: the two objectives coincide for linear preferences without the matching constraint (Vazirani, 2007), but with the matching constraint the concepts are not equivalent. Both NSW and CEEI outcomes are Pareto efficient, but to our knowledge, in matching markets, the agents’ utilities under the two criteria have not been directly compared. Contrasting with CEEI, for stochastic matchings, the NSW outcomes can be calculated by a convex program, i.e., a program that optimizes the product of utilities over the marginal probabilities given by a doubly-stochastic matrix, and is therefore computationally tractable.

A second line of literature considers ordinal mechanisms for one-sided matching. The *random serial dictatorship* (RSD) mechanism has a long history of practical application. Recently it has been used in applications such as housing and course allocation. Pathak and Sethuraman (2011) study the use of RSD for school choice in New York City. RSD is truthful, ex post Pareto efficient, and easy to implement (e.g., Abdulkadiroglu and Sonmez, 1998). On the other hand, RSD is neither ex ante Pareto efficient nor envy-free. To remedy this deficiency of RSD, Bogomolnaia and Moulin (2001) developed the *probabilistic serial* (PS) mechanism which, while not truthful, is ordinally efficient, envy-free, and easy to implement. PS has been studied in various contexts ranging from school assignments to kidney matching and it is often contrasted with RSD. For example, Pathak and Sethuraman (2011) show that students often obtain a more desirable random assignment from PS than from RSD. Nonetheless, under a large market assumption PS and RSD converge and the desirable properties of both are attained (Kojima and Manea, 2010; Che and Kojima, 2010). More recent work has also further studied and compared the performance of these two mechanisms with respect to different metrics both theoretically and experimentally (e.g., Aziz et al. (2016); Hosseini et al. (2018)).

Several recent papers have considered approximation in one-sided matching markets without money when agents have cardinal preferences. With cardinal preferences, it is possible to consider the aggregate welfare of an allocation as the sum of the expected utilities of each agent. For an aggregate notion of welfare to make sense, the values of the agents need to be normalized. Two common normalizations are unit-sum, which scales each agent’s values so that their sum is

<table style="width: 100%; border-collapse: collapse;"> <tr> <th style="width: 33%;"></th> <th style="width: 33%; text-align: center;"><i>A</i></th> <th style="width: 33%; text-align: center;"><i>B</i></th> <th style="width: 33%; text-align: center;"><i>C</i></th> </tr> <tr> <th style="text-align: left;"><i>a</i></th> <td style="text-align: center;">1</td> <td style="text-align: center;">2</td> <td style="text-align: center;">0</td> </tr> <tr> <th style="text-align: left;"><i>b</i></th> <td style="text-align: center;">0</td> <td style="text-align: center;">2</td> <td style="text-align: center;">1</td> </tr> <tr> <th style="text-align: left;"><i>c</i></th> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> </tr> </table>		<i>A</i>	<i>B</i>	<i>C</i>	<i>a</i>	1	2	0	<i>b</i>	0	2	1	<i>c</i>	0	0	1	<table style="width: 100%; border-collapse: collapse;"> <tr> <th style="width: 33%;"></th> <th style="width: 33%; text-align: center;"><i>A</i></th> <th style="width: 33%; text-align: center;"><i>B</i></th> <th style="width: 33%; text-align: center;"><i>C</i></th> <th style="width: 33%; text-align: center;"><i>q</i></th> </tr> <tr> <th style="text-align: left;"><i>a</i></th> <td style="text-align: center;">1</td> <td style="text-align: center;">2</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> </tr> <tr> <th style="text-align: left;"><i>b</i></th> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> <td style="text-align: center;">1/2</td> <td style="text-align: center;">0</td> </tr> <tr> <th style="text-align: left;"><i>c</i></th> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> <td style="text-align: center;">1</td> <td style="text-align: center;">0</td> </tr> </table>		<i>A</i>	<i>B</i>	<i>C</i>	<i>q</i>	<i>a</i>	1	2	0	1	<i>b</i>	0	1	1/2	0	<i>c</i>	0	0	1	0	<table style="width: 100%; border-collapse: collapse;"> <tr> <th style="width: 33%;"></th> <th style="width: 33%; text-align: center;"><i>A</i></th> <th style="width: 33%; text-align: center;"><i>B</i></th> <th style="width: 33%; text-align: center;"><i>C</i></th> <th style="width: 33%; text-align: center;"><i>q</i></th> </tr> <tr> <th style="text-align: left;"><i>a</i></th> <td style="text-align: center;">2/3</td> <td style="text-align: center;">4/3</td> <td style="text-align: center;">0</td> <td style="text-align: center;">0</td> </tr> <tr> <th style="text-align: left;"><i>b</i></th> <td style="text-align: center;">0</td> <td style="text-align: center;">4/3</td> <td style="text-align: center;">2/3</td> <td style="text-align: center;">0</td> </tr> </table>		<i>A</i>	<i>B</i>	<i>C</i>	<i>q</i>	<i>a</i>	2/3	4/3	0	0	<i>b</i>	0	4/3	2/3	0
	<i>A</i>	<i>B</i>	<i>C</i>																																																		
<i>a</i>	1	2	0																																																		
<i>b</i>	0	2	1																																																		
<i>c</i>	0	0	1																																																		
	<i>A</i>	<i>B</i>	<i>C</i>	<i>q</i>																																																	
<i>a</i>	1	2	0	1																																																	
<i>b</i>	0	1	1/2	0																																																	
<i>c</i>	0	0	1	0																																																	
	<i>A</i>	<i>B</i>	<i>C</i>	<i>q</i>																																																	
<i>a</i>	2/3	4/3	0	0																																																	
<i>b</i>	0	4/3	2/3	0																																																	
<i>t</i>	0	1	1	<i>t</i>	2/3	4/3	2/3																																														
i. Agent valuations	ii. Initial solution			iii. Final solution																																																	

Table 2: Simple instance involving three agents; a, b, c ; and three items; A, B, C . The value of agent b in the Nash bargaining solution after the removal of agent c drops by a factor of $4/3$. Normalized values are depicted in (ii) and (iii) along with prices t and q ; these values are depicted in bold-face if the allocation probability of the NSW solution is non-zero.

one, and unit-range, which scales and shifts each agent’s values so that the minimum value is zero and the maximum value is one. Under either of these normalizations, Filos-Ratsikas et al. (2014) show that randomized serial dictatorship is an $\Theta(\sqrt{n})$ approximation and that no algorithm for mapping ordinal preferences to allocations is asymptotically better. Christodoulou et al. (2016) consider the unit-sum normalization and show that the price of anarchy of PS is $\Theta(\sqrt{n})$ and that no mechanism, ordinal or cardinal, is asymptotically better. Important comparison of these above results to ours are as follows: our guarantees do not require a normalization of values. Our approximation guarantees are on per-agent utilities, not on the aggregate welfare which allows some agents to be harmed if other agents benefit. We show that our randomized partial improvement mechanism is asymptotically better than RSD in our per-agent analysis framework by a factor of $\Omega(\sqrt{n})$.

More recently, Immorlica et al. (2017) use a notion of approximate Pareto efficiency to analyze the *raffles mechanism* in one-sided matching markets. This efficiency measure provides per-agent approximation guarantees with respect to the Pareto frontier. Our approximation measure can therefore be thought of as a refinement where instead we compare the agent utilities to a specific highly desired point on the Pareto frontier (the Nash bargaining solution). Instead of eliciting the preferences of the agents, the raffles mechanism instead provides the agents with tickets that they can allocate to items, and items are distributed in proportion to the allocated tickets. As a result, this mechanism is not truthful, but the main result shows that its Nash equilibria are $e/(e-1)$ -approximately Pareto efficient, i.e., that there is no equilibrium where each agent’s utility is increased by more than an $e/(e-1)$ factor.

Our mechanism is based on the *partial allocation* (PA) mechanism of Cole et al. (2013) that truthfully and approximately solves the fair division of heteroge-

neous goods. A novel feature of the PA mechanism is that a fraction of the fair allocation is withheld from individual agents in a way that behaves, in the agents’ utilities, as payments that align the incentives of the agents with the Nash social welfare objective. The fair division problem is closely tied to the cake-cutting literature, which originated in the social sciences but has garnered interest from computer scientists and mathematicians alike (Brams and Taylor, 1996; Moulin, 2003; Robertson and Webb, 1998; Young, 1995). The cake – a heterogeneous, divisible item – is represented by the interval $[0, 1]$ and the agents have valuation functions assigning each subinterval to a non-negative value. These valuations are also assumed to be additive. Algorithmic challenges in cake cutting have recently attracted the attention of computer scientists. A historical overview as well as notable results in cake cutting can be found in surveys by Procaccia (2013) and Procaccia and Moulin (2016). The cardinal matching problem we consider is closely related to the cake cutting problem with piecewise uniform valuations since our agents have linear preferences over items.

Random sampling techniques are now common in the literature on mechanism design. They have been primarily developed for revenue maximization problems where the seller lacks prior information on the agents’ preferences (Hartline and Karlin, 2007). Our use of random sampling more closely resembles the literature on redistribution mechanisms, where the designer aims to maximize the consumer surplus and monetary transfers between agents are allowed (Cavallo, 2006; Guo and Conitzer, 2007). An approach by Moulin (2009) is to single out a random agent as the residual claimant, run an efficient mechanism on the remaining agents, and pay the revenue generated by the mechanism to the residual claimant. Similarly, our mechanism randomly partitions the agents into two groups and attempts to implement the PA mechanism on the first group while using the

items that would be reserved for the second group to implement the first group’s outside option. Further connections between our approach and redistribution mechanisms may be possible.

8 Conclusion and Future Work

We defined the random partial improvement (RPI) mechanism for one-sided matching markets without monetary transfers. RPI both truthfully elicits the cardinal preferences of the agents and outputs a distribution over matchings that approximates every agent’s utility in the Nash bargaining solution.

Our analysis suggests several open questions and directions for future work. A natural open question is whether there exists a truthful mechanism that can achieve a constant factor approximation of the Nash bargaining benchmark. The main obstacle for the RPI mechanism was the non-monotonicity of the Nash bargaining benchmark, so it would be interesting to see if some other mechanism could circumvent this issue. Alternatively, since the construction leading to the lower bound is quite artificial, are there any natural assumptions regarding the valuations of the agents that would mitigate the non-monotonicity?

Another interesting direction would be to study how the utilities of agents in the CEEI outcomes compare to those of the Nash bargaining solution. Recall that the CEEI and the Nash bargaining solution are equivalent in linear markets without the matching constraint (Vazirani, 2007), but are different for matching markets.

Finally, our paper provides a non-trivial mechanism aiming to approximate a well-motivated ex-ante Pareto efficient outcome. One could also consider the design of truthful mechanisms aiming to approximate alternative benchmarks on the ex-ante Pareto frontier. Natural candidates would be the utilitarian (or the egalitarian) outcome which maximize the sum (or the minimum) of the agents’ utilities. One drawback of these outcomes is that, unlike the Nash bargaining solution, they are not scale invariant, but one could consider scaled variants of their objectives, e.g., where the agent values are normalized so that $\sum_{j \in M} v_{i,j} = 1$ for every agent i .

References

- Abdulkadiroglu, A. and Sonmez, T. (1998). Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66:689–701.
- Alaei, S., Jalaly Khalilabadi, P., and Tardos, E. (2017). Computing equilibrium in matching markets. In *ACM Conference on Economics and Computation, EC ’17*, pages 245–261. ACM.
- Aziz, H., Filos-Ratsikas, A., Chen, J., Mackenzie, S., and Mattei, N. (2016). Egalitarianism of random assignment mechanisms: (extended abstract). In *International Conference on Autonomous Agents and Multiagent Systems*, pages 1267–1268.
- Barman, S., Krishnamurthy, S. K., and Vaish, R. (2018). Finding fair and efficient allocations. In *ACM Conference on Economics and Computation, EC ’18*, pages 557–574.
- Bogomolnaia, A. and Moulin, H. (2001). A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295 – 328.
- Brams, S. and Taylor, A. (1996). *Fair Division: From Cake Cutting to Dispute Resolution*.
- Brânzei, S., Gkatzelis, V., and Mehta, R. (2017). Nash social welfare approximation for strategic agents. In *ACM Conference on Economics and Computation, EC ’17*, pages 611–628.
- Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103.
- Budish, E., Cachon, G. P., Kessler, J. B., and Othman, A. (2016). Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, 65(2):314–336.
- Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., and Wang, J. (2016). The unreasonable fairness of maximum Nash welfare. In *ACM Conference on Economics and Computation, EC ’16*, pages 305–322.
- Cavallo, R. (2006). Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 882–889. ACM.
- Che, Y.-K. and Kojima, F. (2010). Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica*, 78(5):1625–1672.
- Christodoulou, G., Filos-Ratsikas, A., Frederiksen, S. K. S., Goldberg, P. W., Zhang, J., and Zhang, J. (2016). Social welfare in one-sided matching mechanisms. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 30–50.

- Cole, R. and Gkatzelis, V. (2018). Approximating the Nash social welfare with indivisible items. *SIAM J. Comput.*, 47(3):1211–1236.
- Cole, R., Gkatzelis, V., and Goel, G. (2013). Mechanism design for fair division: Allocating divisible items without payments. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC '13*, pages 251–268. ACM.
- Devanur, N. R. and Kannan, R. (2008). Market equilibria in polynomial time for fixed number of goods or agents. In *IEEE Symposium on Foundations of Computer Science, 2008. FOCS'08*, pages 45–53.
- Filos-Ratsikas, A., Frederiksen, S. K. S., and Zhang, J. (2014). Social welfare in one-sided matchings: Random priority and beyond. In *International Symposium on Algorithmic Game Theory*, pages 1–12. Springer.
- Garg, J., Hoefer, M., and Mehlhorn, K. (2018). Approximating the Nash social welfare with budget-additive valuations. In *ACM-SIAM Symposium on Discrete Algorithms, SODA '18*, pages 2326–2340.
- Goldberg, A. V., Hartline, J. D., Karlin, A. R., Saks, M., and Wright, A. (2006). Competitive auctions. *Games and Economic Behavior*, 55(2):242–269.
- Guo, M. and Conitzer, V. (2007). Worst-case optimal redistribution of VCG payments. In *ACM Conference on Economics and Computation, EC '07*, pages 30–39. ACM.
- Hartline, J. D. and Karlin, A. R. (2007). Profit maximization in mechanism design. In Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V., editors, *Algorithmic game theory*, chapter 13. Cambridge University Press Cambridge.
- Hartline, J. D. and Roughgarden, T. (2008). Optimal mechanism design and money burning. In *ACM Symposium on Theory of Computing, STOC '08*, pages 75–84.
- Hosseini, H., Larson, K., and Cohen, R. (2018). Investigating the characteristics of one-sided matching mechanisms under various preferences and risk attitudes. *Autonomous Agents and Multi-Agent Systems*, 32(4):534–567.
- Hylland, A. and Zeckhauser, R. (1979). The efficient allocation of individuals to positions. *Journal of Political economy*, 87(2):293–314.
- Immorlica, N., Lucier, B., Weyl, G., and Mollner, J. (2017). Approximate efficiency in matching markets. In *International Conference on Web and Internet Economics*, pages 252–265. Springer.
- Kojima, F. and Manea, M. (2010). Incentives in the probabilistic serial mechanism. *Journal of Economic Theory*, 145(1):106 – 123.
- Moulin, H. (2003). *Fair Division and Collective Welfare*.
- Moulin, H. (2009). Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic theory*, 144(1):96–119.
- Nash, J. (1950). The bargaining problem. *Econometrica*, 18(2):155–162.
- Othman, A., Papadimitriou, C., and Rubinstein, A. (2016). The complexity of fairness through equilibrium. *ACM Transactions on Economics and Computation (TEAC)*, 4(4):20.
- Othman, A., Sandholm, T., and Budish, E. (2010). Finding approximate competitive equilibria: Efficient and fair course allocation. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 873–880.
- Pathak, P. A. and Sethuraman, J. (2011). Lotteries in student assignment: An equivalence result. *Theoretical Economics*, 6(1):1–17.
- Procaccia, A. D. (2013). Cake cutting: Not just child’s play. *Communications of the ACM*, 56(7):78–87.
- Procaccia, A. D. and Moulin, H. (2016). *Cake Cutting Algorithms*, page 311–330. Cambridge University Press.
- Robertson, J. and Webb, W. (1998). *Cake Cutting Algorithms: Be Fair If You Can*.
- Thomson, W. (1983). The Fair Division of a Fixed Supply Among a Growing Population. *Mathematics of Operations Research*, 8(3):319–326.
- Thomson, W. and Lensberg, T. (1989). *Axiomatic Theory of Bargaining with a Variable Number of Agents*. Cambridge University Press.
- Vazirani, V. V. (2007). Combinatorial algorithms for market equilibrium. In Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V., editors, *Algorithmic game theory*, chapter 5, pages 103–134. Cambridge University Press Cambridge.
- Young, H. P. (1995). *Equity: In Theory and Practice*.

Zhou, L. (1990). On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory*, 52(1):123–135.

A The Lower Bound Construction

The construction uses a collection of overlapping submarkets named M_0, M_1, \dots, M_s . Associated with these markets are integer parameters k_r , for $1 \leq r \leq s$. There will be k_r copies of M_r in the construction. As we shall see, $k_s = 1$, thus there will be exactly one copy of M_s .

Next, in Figure 1 and Table 3, we show the form of M_0 . The nodes in this figure correspond to the items, and a directed edge (α, β) , labeled by the name of an agent, indicates that this agent was allocated portions of item α in the initial solution and portions of β in the final solution. Items and bidders occur with multiplicity possibly greater than 1 and this is called their size. In the initial equilibrium, every item is fully allocated as the total size of the bidders and items are the same; in the final equilibrium, item A_0 is the one item that is not fully allocated. Note that in each equilibrium, the conditions (6.3)–(6.5) from Section 6.2 are satisfied.

We continue by presenting the constructions of markets M_r , for $1 \leq r < s$ and of market M_s in Tables 4 and 5, respectively. M_s is very similar to M_r ; the only difference lies in the presence of one additional item I_s , which is the item e_s , the losing bidder, will receive in the final equilibrium, plus one additional bidder, i_s , who leaves in the final setting. In these markets, all items are fully allocated in both equilibria.

To complete the construction we have to show that the various unspecified parameters can be chosen so that the conditions of (6.5) are satisfied for every item-bidder pair. (It is immediate that (6.3)–(6.4) are satisfied.)

LEMMA A.1. *There are choices of values for the unspecified parameters for which the valuations specified above yield the claimed initial and final equilibria.*

Proof. We need to choose the values $v_{f,r}^I, v_{g,r}^I, v_{h,r}^I, v_{f,r}^F, v_{g,r}^F, v_{h,r}^F$, the sizes k_r , and the size proportionality factors $s_{a,r}, s_{b,r}, s_{c,r}, s_{d,r}, s_{f,r}, s_{g,r}, s_{h,r}$ so that for each buyer in each equilibrium, its average value is 1, for $1 \leq r \leq s$. First, we set $s_{a,r} = s_{b,r} + s_{f,r} + s_{c,r}$, $s_{c,r} = s_{d,r} + s_{g,r}$, and $s_{b,r} + s_{d,r} = 1 + s_{h,r}$.

Next, we observe that because the buyer $h_0 = a_1$, its values for $H_0 = A_1$ are the same, i.e. $2 = \frac{8}{9}v_1$, or $v_1 = 2 \cdot \frac{9}{8}$. Similarly, item $E_r = A_{r+1}$, so $v_{r+1} = \frac{9}{8}v_r$. We conclude that $v_r = 2 \cdot \left(\frac{9}{8}\right)^r$.

Now, for buyer b_r , we choose $s_{b,r} = \frac{5}{9}s_{d,r}$, for

$$\frac{\frac{16}{7}s_{b,r}k_r + \frac{2}{7}s_{d,r}k_r}{s_{b,r}k_r + s_{d,r}k_r} = \frac{\frac{16}{7} \cdot \frac{5}{9} + \frac{2}{7}}{\frac{5}{9} + 1} = 1.$$

Thus $s_{d,r} = \frac{9}{14}(1 + s_{h,r})$ and $s_{b,r} = \frac{5}{14}(1 + s_{h,r})$. To ensure these values are integers, we will make sure that $1 + s_{h,r}$ is an integer multiple of 14.

We turn to the values $v_{f,r}^F, v_{g,r}^F, v_{h,r}^F$. We choose $s_{h,r} = 14\lfloor v_r/14 \rfloor + 13$, and $v_{h,r}^F$ to satisfy $v_{h,r}^F s_{h,r} + v_r = s_{h,r} + 1$; i.e. $v_{h,r}^F = (\lfloor v_r \rfloor + 1 - v_r)/\lfloor v_r \rfloor$. We need to confirm that $v_{h,r}^F \leq \frac{9}{7}$; but $v_{h,r}^F = 9/(7v_r) \cdot v_{h,r}^F < v_{h,r}^F < 1$, as $v_r \geq \frac{9}{4}$.

Similarly, when $v_r > \frac{9}{2}$, we set $s_{g,r} = \lfloor \frac{2}{9}v_r \rfloor s_{d,r}$ (for when $v_r \leq \frac{9}{2}$, this would set $s_{g,r} = 0$), and $v_{g,r}^F = s_{g,r} + \frac{2}{9}v_r s_{d,r} = s_{g,r} + s_{d,r}$; i.e. $v_{g,r}^F = (\lfloor \frac{2}{9}v_r \rfloor s_{d,r} + s_{d,r} - \frac{2}{9}v_r s_{d,r})/\lfloor \frac{2}{9}v_r \rfloor s_{d,r}$. Again, we need to confirm that $v_{g,r}^F \leq \frac{3}{2}$; but $v_{g,r}^F = 9/(4v_r) \cdot v_{g,r}^F \leq v_{g,r}^F \leq 1$, as $v_r > \frac{9}{2}$.

When $v_r < \frac{9}{2}$, we set $s_{g,r} = \lceil \frac{9-2v_r}{\frac{2}{3}v_r-1} \rceil$ and $v_{g,r}^F = (s_{g,r} + s_{d,r} - \frac{2}{9}v_r s_{d,r})/s_{g,r}$; but then $s_{d,r} = 9$, so $v_{g,r}^F = 1 + (9 - 2v_r)/\lceil \frac{9-2v_r}{\frac{2}{3}v_r-1} \rceil$. Again, we need to confirm that $v_{g,r}^F \leq \frac{3}{2}$; but $v_{g,r}^F = 9/(4v_r) \cdot v_{g,r}^F \leq 9/(4v_r)(1 + \frac{2}{3}v_r - 1) = \frac{3}{2}$.

As $v_r = 2 \cdot \left(\frac{9}{8}\right)^r$, $v_r \neq \frac{9}{2}$ for any r .

Also, we set $s_{f,r} = \lfloor \frac{16}{9}v_r \rfloor s_{b,r} + \lfloor \frac{4}{9}v_r \rfloor s_{c,r}$ and $v_{f,r}^F s_{f,r} + \frac{16}{9}v_r s_{b,r} + \frac{4}{9}v_r s_{c,r} = s_{f,r} + s_{b,r} + s_{c,r}$; i.e. $v_{f,r}^F = (\lfloor \frac{16}{9}v_r \rfloor s_{b,r} + \lfloor \frac{4}{9}v_r \rfloor s_{c,r} + s_{b,r} + s_{c,r} - (\frac{16}{9}v_r s_{b,r} + \frac{4}{9}v_r s_{c,r}))/(\lfloor \frac{16}{9}v_r \rfloor s_{b,r} + \lfloor \frac{4}{9}v_r \rfloor s_{c,r})$.

We can now calculate the following values.

$$s_{h,r} = 14\lfloor v_r/14 \rfloor + 13$$

$$s_{d,r} = 9(\lfloor v_r/14 \rfloor + 1)$$

$$s_{b,r} = 5(\lfloor v_r/14 \rfloor + 1)$$

$$s_{g,r} = 9\lfloor 2v_r/9 \rfloor (\lfloor v_r/14 \rfloor + 1) \text{ for } v_r > \frac{9}{2}$$

$$s_{g,r} = \lceil \frac{9-2v_r}{\frac{2}{3}v_r-1} \rceil \text{ for } v_r < \frac{9}{2}$$

$$s_{c,r} = 9(\lfloor v_r/14 \rfloor + 1) \cdot (\lfloor 2v_r/9 \rfloor + 1) \text{ for } v_r > \frac{9}{2}$$

$$s_{c,r} = 9 + \lceil \frac{9-2v_r}{\frac{2}{3}v_r-1} \rceil \text{ for } v_r < \frac{9}{2}$$

$$s_{f,r} = 5\lfloor 16v_r/9 \rfloor (\lfloor v_r/14 \rfloor + 1)$$

$$+ 9\lfloor 4v_r/9 \rfloor (\lfloor v_r/14 \rfloor + 1) \cdot (\lfloor 2v_r/9 \rfloor + 1) \text{ for } v_r > \frac{9}{2}$$

$$s_{f,r} = 5\lfloor 16v_r/9 \rfloor (\lfloor v_r/14 \rfloor + 1) \text{ for } v_r < \frac{9}{2}$$

$$s_{a,r} = 5(\lfloor v_r/14 \rfloor + 1) + 9(\lfloor v_r/14 \rfloor + 1) \cdot (\lfloor 2v_r/9 \rfloor + 1)$$

$$+ 5\lfloor 16v_r/9 \rfloor (\lfloor v_r/14 \rfloor + 1)$$

$$+ 9\lfloor 4v_r/9 \rfloor (\lfloor v_r/14 \rfloor + 1) \cdot (\lfloor 2v_r/9 \rfloor + 1) \text{ for } v_r > \frac{9}{2}$$

$$s_{a,r} = 14 + \lceil \frac{9-2v_r}{\frac{2}{3}v_r-1} \rceil + 5\lfloor 16v_r/9 \rfloor \text{ for } v_r < \frac{9}{2}.$$

We also set $k_{r-1} = s_{a,r}k_r$ for $0 \leq r < s$, and create k_0

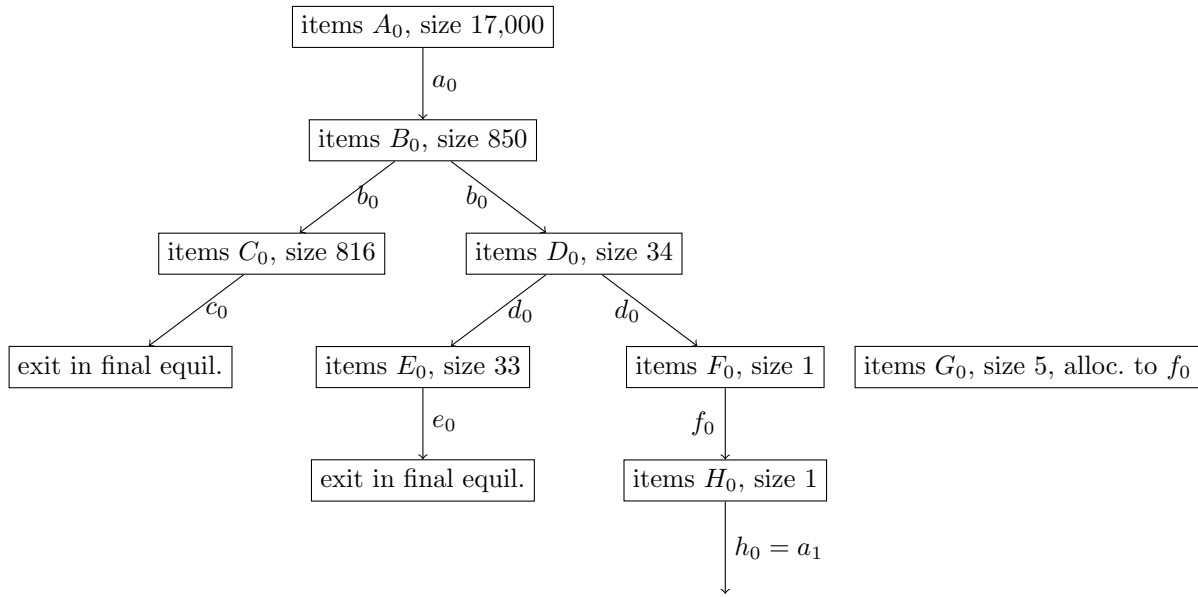


Figure 1: The Allocations in Market M_0 .

item size	A_0 17,000	B_0 850	C_0 816	D_0 34	E_0 33	F_0 1	G_0 5	$H_0 = A_1$ 1	q value
The initial equilibrium:									
Bidder	size								0
a_0	17,000	1	1.5	0	0	0	0	0	1
b_0	850	0	1	$\frac{4687}{7008}$	1.5	0	0	0	0.5
c_0	816	0	0	1	0	0	0	0	0
d_0	34	0	0	0	1	$\frac{5}{11}$	2	0	0
e_0	33	0	0	0	0	1	0	0	0
f_0	6	0	0	0	0	0	$\frac{8}{3}$	$\frac{2}{3}$	$\frac{2}{3}$
h_0	1	0	0	0	0	0	0	1	0
t value		0	0.5	1	1	2	0	1	
The final equilibrium:									
a_0	17,000	$\frac{40}{41}$	$\frac{60}{41}$	0	0	0	0	0	$\frac{40}{41}$
b_0	850	0	$\frac{292}{205}$	$\frac{4687}{4920}$	$\frac{438}{205}$	0	0	0	$\frac{192}{205}$
d_0	34	0	0	0	2	$\frac{10}{11}$	4	0	$\frac{4}{5}$
f_0	6	0	0	0	0	0	$\frac{16}{5}$	$\frac{4}{5}$	2
$h_0 = a_1$	1	0	0	0	0	0	0	0	2
t value		0	$\frac{20}{41}$	$\frac{79}{4920}$	$\frac{6}{5}$	$\frac{6}{55}$	$\frac{16}{5}$	$\frac{4}{5}$	$2 = t_1^{A,F}$

Table 3: Market M_0 , showing normalized valuations, multiplicity of bidders and items (their sizes), assignments (in bold), and the t and q values, for both the initial and final equilibria. The overlap with Market M_1 lies in item H_0 which is also item A_1 and bidder h_0 who is also bidder a_1 . Note that the t values for $H_0 = A_1$ are the same in markets M_0 and M_1 in both equilibria, as are the q values for $h_0 = a_1$.

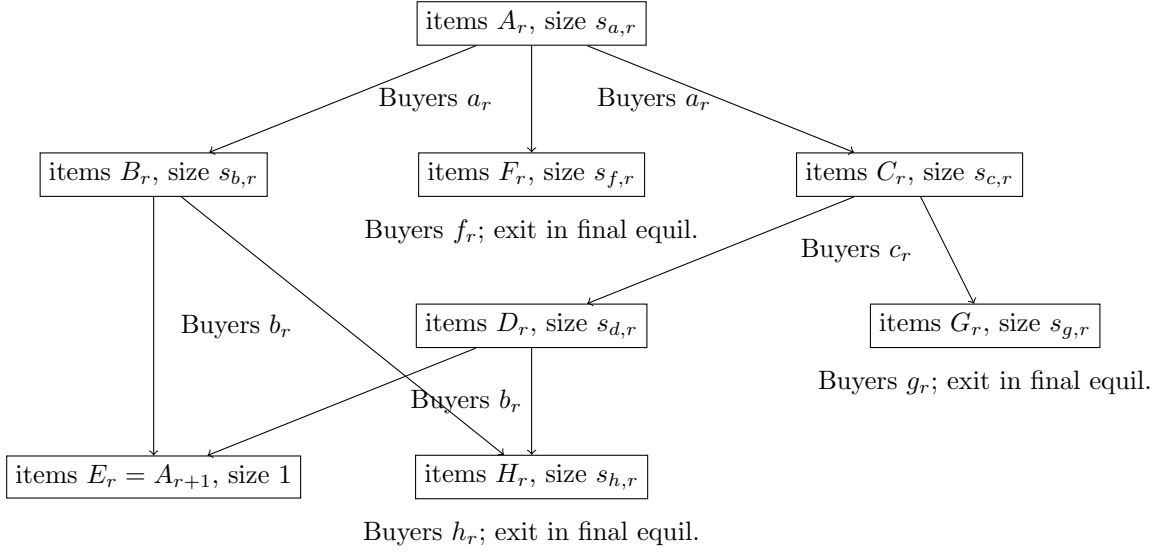



Figure 2: The Allocations in Market M_r , $1 \leq r < s$. The parameters are specified in the proof of Lemma A.1.

item		A_r	B_r	C_r	D_r	E_r $= A_{r+1}$	F_r	G_r	H_r	q value
size		$s_{a,r}$	$s_{b,r}$	$s_{c,r}$	$s_{d,r}$	1	$s_{f,r}$	$s_{g,r}$	$s_{h,r}$	
Initial equilibrium:										
Bidder	size									
a_r	$s_{a,r}$	1	2	$\frac{1}{2}$	0	0	$v_{f,r}^I$	0	0	0
b_r	$s_{b,r} + s_{d,r}$	0	$\frac{16}{7}$	0	$\frac{2}{7}$	$\frac{9}{7}$	0	0	$v_{h,r}^I$	$\frac{2}{7}$
c_r	$s_{c,r}$	0	0	1	$\frac{1}{2}$	0	0	$v_{g,r}^I$	0	$\frac{1}{2}$
f_r	$s_{f,r}$	0	0	0	0	0	1	0	0	0
g_r	$s_{g,r}$	0	0	0	0	0	0	1	0	0
h_r	$s_{h,r}$	0	0	0	0	0	0	0	1	0
t value		1	2	$\frac{1}{2}$	0	1	1	1	1	
Final equilibrium:										
a_r	$s_{a,r}$	$\frac{8}{9}v_r$	$\frac{16}{9}v_r$	$\frac{4}{9}v_r$	0	0	$v_{f,r}^F$	0	0	0
b_r	$s_{b,r} + s_{d,r}$	0	$\frac{16}{9}v_r$	0	$\frac{2}{9}v_r$	v_r	0	0	$v_{h,r}^F$	0
c_r	$s_{c,r}$	0	0	$\frac{4}{9}v_r$	$\frac{2}{9}v_r$	0	0	$v_{g,r}^F$	0	0
t value		$\frac{8}{9}v_r$	$\frac{16}{9}v_r$	$\frac{4}{9}v_r$	$\frac{2}{9}v_r$	v_r	$v_{f,r}^F$	$v_{g,r}^F$	$v_{h,r}^F$	

Table 4: Component C_r , showing normalized valuations, multiplicity of bidders and items (their sizes), and the t and q values, for both the initial and final equilibria. $v_r^{a,F}$ and $v_r^{c,F}$ are normalizing factors, equal to the value of the assignments using the initial valuations.



Buyer i_s ; exit in final equil.

Figure 3: The Allocations in Market M_s .

item size	E_s 1	I_s 1	q value
The initial equilibrium:			
Bidder	size		
e_s	1	1	$1/(v_s + 1)$
i_s	1	0	1
t value	1	1	
The final equilibrium:			
e_s	1	$v_s + 1$	1
t value	v_s	0	

Table 5: Component C_s , showing the additional portion in addition to the part shown in Table 4.

copies of M_0 . Recall that $k_s = 1$. \square

To conclude the lower bound analysis we lower bound the size of s and hence of v_s . We observe that for $v_r > \frac{9}{2}$, $s_{a,r} \leq 5(v_r/14 + 1) + 9(v_r/14 + 1) \cdot (2v_r/9 + 1) + 5 \cdot 16v_r/9(v_r/14 + 1) + 4v_r \cdot (v_r/14 + 1) \cdot (2v_r/9 + 1) \leq 4v_r^3/63 + 91v_r^2/63 + 143v_r/9 + 14 \leq 2v_r^3$, as $v_r \geq \frac{9}{2}$, and for $v_r < \frac{9}{2}$, $s_{a,r} \leq 14 + 9 + 80v_r/9 \leq 4v_r^3$.

Note that $v_s = 2 \left(\frac{9}{8}\right)^s$. We can conclude that

$$\begin{aligned}
k_0 &\leq 4v_s^3 \cdot 4 \left(\frac{8}{9}v_s\right)^3 \dots 4 \left(\left(\frac{8}{9}\right)^{s-1} v_s\right)^3 \\
&= 4 \left(2 \left(\frac{9}{8}\right)^s\right)^3 \cdot 4 \left(2 \left(\frac{9}{8}\right)^{s-1}\right)^3 \dots \left(2 \left(\frac{9}{8}\right)\right)^3 \\
&= (32)^s \left(\frac{9}{8}\right)^{s(s-1)/2} \\
&\leq \left(\frac{9}{8}\right)^{(s^2+58s)/2}.
\end{aligned}$$

Thus the size of item A_0 is at most $17000 \left(\frac{9}{8}\right)^{(s^2+58s)/2}$. But this is more than $n/2$ by inspection of the construction. Therefore we can choose s in our construction to

satisfy $17000 \left(\frac{9}{8}\right)^{(s^2+58s)/2} \geq n/2$; thus

$$s^2 + 58s)/2 \cdot \log \frac{9}{8} + \log 17000 \geq \log n - 1.$$

and hence

$$(s + 29)^2 - 29^2 \geq 2(\log n - 1 - \log 17000)/\log \frac{9}{8}$$

So,

$$s \geq (2[(\log n) - 1 - \log 17000 + 142]/\log \frac{9}{8})^{1/2} - 29.$$

This implies

$$\begin{aligned}
v_s &\geq \left(\frac{9}{8}\right)^{(2 \log n / \log \frac{9}{8})^{1/2} - 29} \\
&\geq 2\sqrt{2 \log \frac{9}{8} \log n} - 29 \log \frac{9}{8} \\
&\geq 2\sqrt{\log n/2 - 5}.
\end{aligned}$$