Journal of Computer and System Sciences 140 (2024) 103491

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

journal homepage: www.elsevier.com/locate/jcss

Bounded incentives in manipulating the probabilistic serial rule ☆

Haoqiang Huang [a], Zihe Wang [b], Zhide Wei [c], Jie Zhang [d,*]

[a] Hong Kong University of Science and Technology, Hong Kong, China
[b] Renmin University of China, China
[c] Peking University, China
[d] University of Bath, UK

A R T I C L E I N F O

Article history:
Received 17 January 2021
Received in revised form 11 March 2023
Accepted 8 October 2023
Available online 28 October 2023

Keywords:
Incentive ratio
Manipulation
Probabilistic serial mechanism
Random assignment
Resource allocation

A B S T R A C T

The Probabilistic Serial mechanism is valued for its fairness and efficiency in addressing the random assignment problem. However, it lacks truthfulness, meaning it works well only when agents' stated preferences match their true ones. Significant utility gains from strategic actions may lead self-interested agents to manipulate the mechanism, undermining its practical adoption. To gauge the potential for manipulation, we explore an extreme scenario where a manipulator has complete knowledge of other agents' reports and unlimited computational resources to find their best strategy. We establish tight incentive ratio bounds of the mechanism. Furthermore, we complement these worst-case guarantees by conducting experiments to assess an agent's average utility gain through manipulation. The findings reveal that the incentive for manipulation is very small. These results offer insights into the mechanism's resilience against strategic manipulation, moving beyond the recognition of its lack of incentive compatibility.

1. Introduction

Resource allocation is a fundamental and widely applicable area within AI and computer science. The random assignment problem is central in the resource allocation area, which has wide applications in allocating workers to shifts, houses to people, dormitories to students, ramp and hangars to airlines, and chores to cleaning staff. In the problem, there are a set of agents and a set of items. The agents participate in a mechanism by reporting their private preferences over the items. The mechanism then assigns items to agents, according to a pre-defined allocation rule.

The random assignment problem was introduced in [27] and has been studied extensively ever since. Over the years, several mechanisms have been investigated, including Probabilistic Serial [10,29,8,9,23,2], Random Priority [1,34,4,17], and Competitive Equilibrium from Equal Incomes (CEEI) [12]. In the indivisible goods setting, the Top Trading Cycles (TTC) method is well-studied and generalized to investigate various problems. In particular, Abdulkadiroglu and Sönmez [1] proposed an adaptation of the TTC method and established an equivalence between the adapted mechanism and Random

☆ A paper with the same title was accepted by the Thirty-Fourth AAAI Conference on Artificial Intelligence as a full paper. The paper contains results when $n \geq m$. This manuscript extends the previous work by covering the case $n < m$ and more extensive experimental results.

* Corresponding author.
E-mail addresses: hhuangbm@connect.ust.hk (H. Huang), wang.zihe@ruc.edu.cn (Z. Wang), zhidewei@pku.edu.cn (Z. Wei), jz2558@bath.ac.uk (J. Zhang).

https://doi.org/10.1016/j.jcss.2023.103491

Priority. Kesten [30] proposed several extensions of these popular mechanisms and presented an equivalence result between these mechanisms in terms of economic efficiency.

There are various desired properties one would like these mechanisms to fulfill. In terms of *efficiency*, an allocation *Pareto improves* another if items can be re-allocated to make at least one person better off without making any other individual worse off. An allocation is *Pareto efficient* if there is no allocation in which Pareto improves it. In terms of *fairness*, an allocation is *envy-free* if no agent prefers another agent's allocation to its own. An allocation is *proportional* if each of the n agents receives at least $1/n$ of the resources by its own subjective valuation. Also, a mechanism is *equal treatment of equals* if agents with precisely the same valuation functions must have the same probabilities of receiving each item; a mechanism is *nonbossy* if an agent cannot change other agents' allocation without changing its own allocation. *Incentive-compatibility* is an important property in understanding the role that strategic behavior can play. A mechanism is *manipulable* if an agent can misreport its preferences and improve its utility. A mechanism that is not manipulable is *incentive-compatible* (a.k.a., *truthful*).

Truthfulness, in a sense, is on top of the properties mentioned above. In other words, take Pareto efficiency, for example, it means that no other allocation would Pareto improve the allocation output by a mechanism with respect to agents' declared preferences. However, there may exist allocations Pareto improve it with respect to agents' actual preferences. Without the truthfulness guarantee, an agent may manipulate in a mechanism and could potentially cause significant welfare loss and unfairness to the agents.

As one may expect, these properties are not compatible, as shown by several impossibility results. A notable trilemma by Zhou [35] states that there exists no mechanism that is truthful, symmetric, and ex-ante Pareto efficient. Bogomolnaia and Moulin [10] shows that there does not exist a truthful, equal treatment of equals, and ordinal efficient mechanism. Given these impossibility results, one reasonable choice for implementation is the Random Priority mechanism. It is easy to implement and incentive-compatible, but lacks efficiency, as all agents may increase their probabilities of receiving more preferred items by implementing the Probabilistic Serial mechanism. Moreover, Probabilistic Serial is *stochastic-dominance envy-free*. However, it is not truthful. In fact, many mechanisms, such as CEEI, is manipulable as well. Ekici and Kesten [22] shows that when agents are not truthful, the outcome of Probabilistic Serial may not satisfy desirable properties related to efficiency and envy-freeness. Hence, researchers look to develop refined analytic and experimental works to answer a basic question: *which mechanism to employ in practical applications*?

There are at least two scenarios that agents may dispense with manipulation in mechanisms that are not incentive-compatible. One scenario is where computational complexity is considered as an obstacle against manipulation. If manipulation is hard to compute, agents will likely behave sincerely. The other scenario is by characterizing the extent to which strategic manipulations can increase utilities. If agents can increase their utility by only a small extent, given that there is an inherent cost for them to collect necessary information from other agents in order to compute the best response strategy, it is more likely they would behave truthfully.

In this paper, we consider the second scenario in the Probabilistic Serial mechanism. We adopt the *incentive ratio* notion [16] [15] [14] to quantify an agent' incentive to deviate from reporting their actual private information. Informally, it is the factor of the largest possible utility gain that an agent achieves by behaving strategically, given that all other agents have their strategies fixed. Our main result is the following theorem.

Theorem. *In the Probabilistic Serial mechanism, when the number of agents is no less than the number of items, no agent is able to unilaterally manipulate and increase its utility to more than 1.5 times of the utility when reported truthfully; when the number of agents is less than the number of items, this ratio is bound by 2. Both of these ratios are tight.*

In understanding the implication of these approximation results on the manipulation incentives, we put in mind of two technical facts. Firstly, the incentive ratio is defined in a *worst-case* sense. Therefore, its bound is the strongest approximation guarantee of the manipulation incentives in all cases. It also means that the incentive ratio bound ignores the likelihood that extreme cases happen. It could be the case that the probability for an agent to attain the 3/2 or 2 times utility is negligible. Indeed, our tight bound examples are rather pathological as shown in Section 4. Secondly, our results build upon the *complete information* and *perfect rationality* assumptions. That is, agents have complete information about other agents' preferences and are able to compute the best response strategy accordingly. If either of these assumptions is missing, the agents' power to manipulate the mechanism would be much smaller than what these bounds imply. In fact, computing the best response strategy is intractable in general [5]. Moreover, Hugh-Jones, Kurino, and Vanberg [26] showed that humans do not manipulate Probabilistic Serial mechanism optimally. Halpern, Pass, and Seeman [28] provided an excellent survey of work using *bounded rationality* in decision theory. Therefore, small constant incentive ratios, in particular, 3/2 and 2 in our results, indicate that the agents' incentive and ability to manipulate the mechanism is reasonably bounded. To support this statement, we further conduct some experiments. The experiments demonstrate that the agents' utility gain by strategic manipulation is rather limited on average, even if there are not many agents and items.

Another work that initiated this study is the Price-of-Anarchy bound presented by Christodoulou et al. [17]. In the paper, the authors showed that Probabilistic Serial achieves a PoA bound of $\Theta(\sqrt{n})$, which suggests a large efficiency loss due to selfish behavior of the agents. Hence, it is desirable to characterize the extent to which agents can benefit from manipulation in Probabilistic Serial from both the worst-case analysis and beyond worst-case analysis sense.

From the technical point of view, characterizing the incentive ratio of a mechanism is challenging. For a start, determining the best response strategy, i.e., a report that maximizes the agent's expected utility, is often intractable. In fact, this task is NP-hard in the context of the Probabilistic Serial mechanism [5], and only local manipulation has been studied experimentally by greedy search [32]. In a nutshell, our proofs circumvent this obstacle by bounding the time span that other agents would take to eat up the items that an agent is interested in while pausing the agent. This period upper bounds the largest utility the agent may receive by manipulation. To derive a tight bound, we need case studies and a fine-grained analysis on the number of agents who are eating specific items.

1.1. Related work

In the presence of incentives, the random assignment problem has been extensively studied in Computer Science and Economics over the years [35,21,31]. We refer the interested reader to surveys [2,33].

One of the focal points is how precisely agents might compute beneficial manipulations. What if it is just too computationally difficult to compute a manipulation [7,6]? Manipulation has been shown to be computationally hard to compute in many voting situations, e.g., [19,18].

For the Probabilistic Serial mechanism, Aziz et al. [5] showed that computing the best response (manipulation) under complete information of the other agents' strategies is NP-hard for Expected Utility maximization, a.k.a., EU-relation, but polynomial-time algorithm exists for the Downward Lexicographic maximization, a.k.a., DL-relation. In addition, they showed that Nash deviations under the Probabilistic Serial mechanism may cycle, but a pure Nash equilibrium is guaranteed to exist. Unfortunately, computing a pure Nash equilibrium is intractable in general. Bogomolnaia and Moulin [11] considered an approximation problem for maximizing the number of items actually assigned in the context of envy-free rather than truthful mechanisms, and for strict preference lists and un-weighted agents. They showed that PS has an approximation ratio of $\frac{e}{e-1}$, which is tight for any envy-free mechanism.

The empirical work by Hosseini, Larson, and Cohen [25] disclosed some results on the manipulability of the Probabilistic Serial mechanism. Their experiments show that the mechanism is almost always manipulable for some combination of agents and items, and the fraction of strongly manipulable profiles goes to one as the ratio of items to agents increases. However, their results do not reveal the extent to which these beneficial manipulations are.

An interesting work by Che and Kojima [13] showed that Random Priority and Probabilistic Serial mechanisms become equivalent when there exist a large number of copies of each item type. Their results imply that, on the one hand, the inefficiency of the Random Priority mechanism becomes small and disappears in the limit, as the economy becomes large; on the other hand, the incentive problem of the Probabilistic Serial mechanism disappears in large economies.

The incentive ratio notion was proposed by [16] [15]. The authors investigated the buyers' incentive to manipulate Fisher markets. They showed that the incentive ratios of Fisher markets, when the buyers are featured with typical constant elasticity of substitution utility functions (including linear, Leontief, and Cobb-Douglas utility functions), are 2 or 1.445.

2. Preliminaries

The random assignment problem[1] consists of n agents and m divisible items o_j, $j \in \{1, 2, \cdots, m\}$. Each agent may prefer one item to another, and this is their private information. A mechanism aggregates individual agents' preferences and allocates the items to agents accordingly. The Probabilistic Serial (PS) mechanism asks the agents to report items in the order from the most to least favorable order, $\succ$. The allocation process of the mechanism simulates a *simultaneous eating algorithm*. That is, agents simultaneously "eat" their most preferred items at a uniform speed, moving onto their next most preferred item whenever an item is fully eaten. We denote the output of the PS mechanism by a matrix $\mathrm{X} = [x_{ij}]_{n\times m}$, where x_{ij} is the amount of item j that agent i has eaten. Obviously, $\sum_i x_{ij} = 1, \forall j$.

Example 1. There are three agents and three items a, b, and c. The agents' preferences and the corresponding allocation are

$$1: b \succ a \succ c \qquad\qquad \mathrm{X} = \begin{matrix} 0 & 3/4 & 1/4 \\ 1/2 & 1/4 & 1/4 \\ 1/2 & 0 & 1/2 \end{matrix}$$
$$2: a \succ b \succ c$$
$$3: a \succ c \succ b$$

In this example, Agent 1 starts with eating item b alone, while Agents 2 and 3 share item a. By the time Agents 2 and 3 finish eating item a, they turn to eat their next favorable items, b and c, respectively. Agents 1 and 2 share the remaining half of the item b. By the time they finish eating b, Agent 1 would like to eat item a. As item a is eaten up, Agent 1 turns to eat the next favorable and available item c. Agent 2 turns to eat item c as well, and Agent 3 has been eating item c after finishing a. The three agents share the remaining $\frac{3}{4}$ of item c. The process is illustrated in Fig. 1 and hence the allocation matrix X.

[1] Many works consider the special case that $n = m$. In this case, the allocation matrix is doubly stochastic. Thus, any probabilistic allocation can be seen as a convex combination of a set of deterministic allocations, and equivalence between a probabilistic allocation of indivisible items and a fractional assignment of divisible items is established. However, in general, n and m are not necessarily equal.

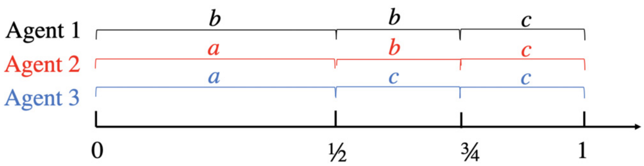

Fig. 1. The simultaneous eating process in Example 1.

We assume that agents have von Neumann-Morgenstern preferences over the items. That is, underlying their ordinal preferences, agents have *cardinal* utilities which accord with their ordinal preferences. The agents use numerical values to specify how much they prefer an item to another. For instance, in Example 1, Agent 1 could have cardinal values 2, 6, and 1 on items a, b, and c, respectively. We denote the utility derived by agent i on obtaining a unit of item j by a_{ij}. Under standard normalization, $0 \le a_{ij} \le 1$, $\forall i \in [n]$, $j \in [m]$. Denote vector $\mathbf{a}_i = (a_{i1}, \cdots, a_{1m})$. So, the expected utility of agent i is $u_i = \sum_j a_{ij} x_{ij}$.

Agents are self-interested and may misreport their ordinal preferences if that results in a better allocation (from their perspective). Denote $\succ_{-i}$ agents' preferences except agent i. Let $\succ_i$ be agent i's true preferences and denote $(\succ_i, \succ_{-i})$ the *truthful profile*; let $\succ_i'$ be a misreport by agent i and denote $(\succ_i', \succ_{-i})$ a *manipulation profile*. Agent i has incentive to report $\succ_i'$ if $u_i(\succ_i, \succ_{-i}) < u_i(\succ_i', \succ_{-i})$.

It is well observed that agents in the PS mechanism may misreport their preference orders [10]. In Example 1, if Agent 1 misreports its preferences as follows, then the allocation changes accordingly.

$$\begin{array}{ll} 1: a \succ b \succ c & \\ 2: a \succ b \succ c & \mathrm{X} = \\ 3: a \succ c \succ b & \end{array} \quad \begin{array}{ccc} 1/3 & 1/2 & 1/6 \\ 1/3 & 1/2 & 1/6 \\ 1/3 & 0 & 2/3 \end{array}$$

For some utility function that is compatible with Agent 1's true preferences $b \succ a \succ c$, for example, $a_{1a} = 0.9, a_{1b} = 1, a_{1c} = 0$, its utilities are $u_1 = 0.9 \times 0 + 1 \times \frac{3}{4} + 0 \times \frac{1}{4} = 0.75$ in the truthful profile and $u_1 = 0.9 \times \frac{1}{3} + 1 \times \frac{1}{2} + 0 \times \frac{1}{6} = 0.8$ in the manipulation profile. Therefore, $u_i(\succ_i, \succ_{-i}) < u_i(\succ_i', \succ_{-i})$ in this scenario. The high-level intuition behind this manipulative example is that, to Agent 1, the difference between the value of items a and b is small, but item a is more competitive as the other two agents place it as their most preferred item. So, starting by eating a more competitive item a followed by item b can prolong the total time of consuming these two items.

The *incentive ratio* characterizes the extent to which utilities can be increased by strategic plays of individuals. The *incentive ratio* of agent i in the PS mechanism is

$$r_i = \sup_{\succ_{-i}} \frac{\sup_{\succ_i'} u_i(\succ_i', \succ_{-i})}{u_i(\succ_i, \succ_{-i})}.$$

Given any report $\succ_{-i}$ of other agents, the denominator is the utility of agent i when it truthfully reports its preferences, and the numerator is the largest possible utility of agent i when it unilaterally misreports its preferences. The incentive ratio of agent i is then the maximum value of the ratio over all possible inputs of other agents and any possible cardinal preferences that it may possess. Therefore, this ratio is defined over all possible preference profiles $(\succ_i, \succ_{-i})$. The incentive ratio of the PS mechanism is the largest ratio amongst all agent, $\max_i r_i$. Throughout the paper, w.l.o.g., we consider the strategic manipulation of Agent 1.

Clearly, the incentive ratio is a worst-case consideration amongst all possible inputs. To characterize the utility gain at the most extreme circumstance, we assume that Agent 1 has complete information about other agents' preferences, $\succ'_{-i}$, and is powerful enough to figure out its best response strategy $\succ_1'$ without any computability constraint. Hence, in Section 3, we quantify the extent to which an agent can benefit from manipulation in the PS mechanism in this extreme setting, regardless of how likely it would happen. In Section 4, we can see that the tight bound examples retain a specific shape and size. In Section 5, the numerical experiments show that the incentive ratio is close to 1 on randomly generated instances.

Throughout the paper, we present the agents' preferences as the strict ordering of items for ease of discussion. In the setting where agents are allowed to be indifferent between items, the incentive ratio bounds in this paper hold unchanged as long as the PS mechanism outputs a deterministic allocation. However, when the PS mechanism allocation is not deterministic on the full preference domain where indifferences and incomplete preference lists are allowed,[2] an agent, evaluated by its real vNM preferences, may receive different utilities over possible allocations when it misreports its preferences. Due to this complication, our techniques and results may not extend to the full preference domain, while PS mechanism allocation is not deterministic.

We present the high-level ideas used in Section 3 to facilitate understanding the upper bound proofs. First, we show that Agent 1's utility gain is maximized when it has either close to 1 or close to 0 values. Therefore, its utility can be

[2] See Example 5.1.3. in [29].

upper-bounded by the time span of eating close-to-1-valued items. Next, the proofs are mostly case-by-case analyses on the extent to which Agent 1 can prolong the eating time of close-to-1-valued items via misreporting. Notably, Agent 1's eating time can be bounded by removing this agent and letting other agents complete his portion under truthful reporting. In order to keep the eating time of other items unaffected, the eating of the rest items needs to be properly paused. With this auxiliary approach, the increase of eating time can be bounded based on the number of agents who eat the portion of the removed agent. A more refined analysis is used when only one agent is eating the portion of the removed agent.

3. Incentive ratio upper bounds

In this section, we prove the incentive ratio upper bounds. The section consists of three parts. In the first part, we establish some preliminary results to facilitate proving the bounds. In the next two parts, we proceed to show the upper bounds for the cases $n \geq m$ and $n < m$, respectively.

3.1. Preliminary results

Our first step is a reduction, which shows that it is sufficient to consider the instances in which Agent 1's cardinal values a_{1j} are either close to 1, or close to 0, $\forall j = 1, \cdots, m$.

Lemma 1. *Given any truthful profile $(\succ_1, \succ_{-1})$ and Agent 1's cardinal preference a_{1j}'s that are compatible with the ordering $\succ_1$, denote ratio $c = \frac{u_1'}{u_1}$, where u_1' is Agent 1's maximum utility attainable by manipulation. Then one can always construct a corresponding preference b_{1j}, where b_{1j} are either close to 1 or close to 0, and are compatible with $\succ_1$, such that the ratio c is no less than before.*

Proof. W.l.o.g., assume $o_1 \succ_1 o_2 \succ_1 \cdots \succ_1 o_m$ and $1 \geq a_{11} > a_{12} > \cdots > a_{1m} \geq 0$. Denote l_j and l_j' the time span that Agent 1 spent on eating item o_j in the truthful profile and the manipulation profile, respectively. By definition,

$$c(a_{1j}) = \frac{u_1'}{u_1} = \frac{a_{11}l_1' + a_{12}l_2' + \cdots + a_{1m}l_m'}{a_{11}l_1 + a_{12}l_2 + \cdots + a_{1m}l_m}.$$

Denote

$$k = \operatorname*{argmax}_j \frac{l_1' + l_2' + \cdots + l_j'}{l_1 + l_2 + \cdots + l_j}, \quad c_{\max} = \frac{\sum_{j=1}^k l_j'}{\sum_{j=1}^k l_j}.$$

Let $a_{1,m+1} = 0$. We can rewrite c as follows, and show that $c(a_{1j}) \leq c_{\max}$.

$$\begin{aligned} c(a_{1j}) &= \frac{\sum_{j=1}^m (a_{1j} - a_{1,j+1}) \sum_{h=1}^j l_h'}{\sum_{j=1}^m (a_{1j} - a_{1,j+1}) \sum_{h=1}^j l_h} \\ &\leq \max_j \left\{ \frac{(a_{1j} - a_{1,j+1}) \sum_{h=1}^j l_h'}{(a_{1j} - a_{1,j+1}) \sum_{h=1}^j l_h} \right\} \\ &= \max_j \left\{ \frac{\sum_{h=1}^j l_h'}{\sum_{h=1}^j l_h} \right\} \\ &= c_{\max}. \end{aligned}$$

We will show that by carefully pushing a_{1j}'s towards binary values 1 or 0, the ratio c is non-decreasing. To this end, we construct a new preference profile $b_{1,j} = 1 - (j-1)\epsilon, j = 1, \cdots, k$ and $b_{1,j} = (m-j)\epsilon, j = k+1, \cdots, m$. Note that $b_{1,j}$'s are compatible with $\succ_1$. So, the truthful allocation l_j's remain the same; by using the same strategy, the manipulation allocation l_j''s are also kept the same. Moreover,

$$\begin{aligned} c(b_{1,j}) &= \frac{\sum_{j=1}^m (b_{1,j} - b_{1,j+1}) \sum_{h=1}^j l_h'}{\sum_{j=1}^m (b_{1,j} - b_{1,j+1}) \sum_{h=1}^j l_h} \\ &= \frac{\epsilon \sum_{j \neq k} \sum_{h=1}^j l_h' + (1 - (m-2)\epsilon) \sum_{h=1}^k l_h'}{\epsilon \sum_{j \neq k} \sum_{h=1}^j l_h + (1 - (m-2)\epsilon) \sum_{h=1}^k l_h} \\ &\to c_{\max} \quad (\epsilon \to 0) \quad \square \end{aligned}$$

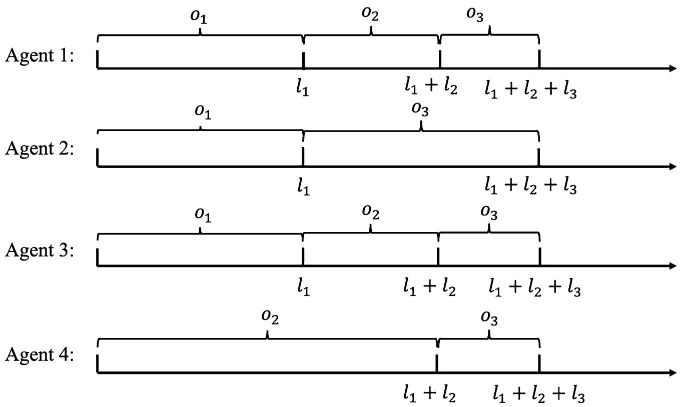

Fig. 2. The normal scenario in Example 2.

The role of this lemma is similar to the *quasi-combinatorial structure* [24] and the *zero-one principle* [3]. It substantiates the intuition that because the mechanism is ordinal, the worst-case incentive ratio is encountered on extreme valuation profiles.

By Lemma 1, we can classify the set of items $O = \{o_1, \cdots, o_m\}$ into two categories. They are, the set of items that Agent 1 is interested in, i.e., $\overline{O} = \{o_j | \ a_{1j} \ \textit{is close to} \ 1\}$, and the set of items that Agent 1 is not interested in, i.e., $\underline{O} = \{o_j | \ a_{1j} \ \textit{is close to} \ 0\}$. For ease of notation, assume that Agent 1 is interested in k items. So, $\overline{O} = \{o_1, \cdots, o_k\}$ and $\underline{O} = \{o_{k+1}, \cdots, o_m\}$. Let x_j be the amount of item j that Agent 1 eats when it reports the true preference $\succ_1$. In this case, its utility $u_1 = \sum_{j=1}^{m} a_{1j} x_j \approx \sum_{j=1}^{k} x_j$. We note that Agent 1 may not be able to get a positive fraction of each item in $\overline{O}$, as some of the items may have been eaten up by other agents when Agent 1 was eating more favorable items. As such, some of these x_j's may be equal to 0, $j = 1, \cdots, k$. Agent 1 may also eat some of the items in $\underline{O}$, but the contribution of these items to Agent 1' utility is negligible.

The agents report their preferences to the PS mechanism and consume items according to the simultaneous eating algorithm. In the eating process, we denote two scenarios to facilitate the following analysis. In the *normal scenario*, all agents eat items according to their reported ordinal preferences as normal; in a *pause scenario*, a set of agents is paused from eating for some time while other agents continue eating normally. Before employing these two scenarios in the following proofs, we present an example for a visual illustration.

Example 2. There are four agents and three items o_1, o_2, o_3 and o_4. The agents' preferences are

1: $o_1 \succ o_2 \succ o_3$
2: $o_1 \succ o_3 \succ o_2$
3: $o_1 \succ o_2 \succ o_3$
4: $o_2 \succ o_3 \succ o_1$

The eating process in the normal scenario is illustrated in Fig. 2. The eating process in a pause scenario, in which Agent 1 is paused for the whole time, Agent 2 is paused between time $\frac{3}{2}l_1 + l_2$ and $\frac{3}{2}(l_1 + l_2)$, and Agent 4 is paused between time l_1 and $\frac{3}{2}l_1$, is illustrated in Fig. 3.

We shall clarify that the set of agents and the period that they are paused when we refer to a pause scenario in the rest of the paper. The following lemma, though, holds for any set of agents and any period. It holds regardless of if the reported preference is a truthful profile or a manipulation profile. It compares the amount of each item that is not eaten up yet in two scenarios.

Lemma 2. *Given agents' reported preferences $(\succ_1, \cdots, \succ_n)$, suppose that the eating process of a set of agents is paused from time t. For any item j, at any moment after time t until it is fully eaten up, the remaining amount of item j in this pause scenario is no less than the remaining amount of item j in the normal scenario.*

Proof. We prove it by contradiction. Denote $t_{\inf}$ the earliest moment, after which there exists an item j^*, its remaining amount in the pause scenario is less than its remaining amount in the normal scenario. For a small enough $t_\delta > 0$, at the moment $t_{\inf} + t_\delta$, the number of agents who are eating item j^* in the pause scenario must be more than the number of agents who are eating item j^* in the normal scenario. Otherwise, the amount of j^* cannot be smaller in the pause scenario. Denote agent i one of these agents who are eating item j^* in the pause scenario but not in the normal scenario. Recall that every agent eats items from the most preferred to the least preferred item according to their reported preference. Then

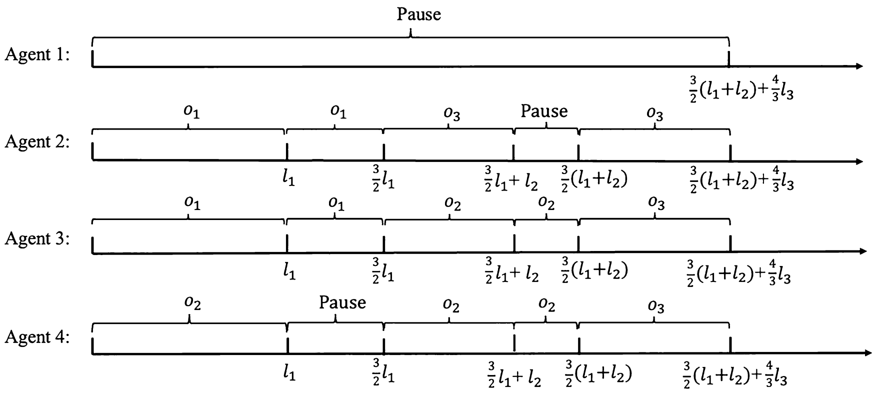

Fig. 3. A pause scenario in Example 2.

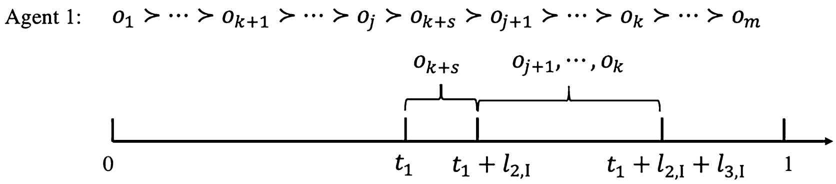

Fig. 4. Agent 1's preference order and the eating process in Scenario I. Given agents' reports, it is the normal scenario without any pause operation.

agent i's eating status implies that there exists an item j', which is more preferable to item j^* in agent i's report, such that at the moment $t_{\text{inf}} + t_\delta$, item j' is eaten up in the pause scenario but is not eaten up yet in the normal scenario. However, it contradicts our choice of item j^* and the definition of t_{inf}. $\square$

We note that Lemma 2 is critical and is repetitively used in the following analysis. It holds regardless of the number of agents that are paused and the time span that they are paused.

Next, we show that the manipulator, Agent 1, would only rearrange the order of items in $\overline{O}$ in improving its utility. In effect, Lemma 3 restricts the searching space when considering profitable manipulations.

Lemma 3. *In any possible misreports that Agent 1 may improve its utility, Agent 1 would not place an item in $\underline{O}$ in front of the items in $\overline{O}$.*

To facilitate the proof of Lemma 3, we first prove a lemma. It compares Agent 1's utility when it places an item in $\underline{O}$ in front of some items in $\overline{O}$ (Scenario I, Fig. 2) and when it places the item after items in $\overline{O}$ (Scenario IV, Fig. 3).

- In Scenario I, Agent 1 reports a preference order in which items $o_{k+1}, \cdots, o_{k+s}$ in $\underline{O}$ are in front of some items in $\overline{O}$. Denote o_{k+s} the last item in $\underline{O}$ that is in front of items in $\overline{O}$. Denote t_1 the moment that Agent 1 starts to eat item o_{k+s}. Denote $l_{2,\text{I}}$ and $l_{3,\text{I}}$ the time span that Agent 1 eats item o_{k+s} and items $o_{j+1}, \cdots, o_k$, respectively. See Fig. 4.
- Scenario II is a pause scenario of Scenario I. Specifically, starting from moment t_1, all agents except Agent 1 are paused from eating, until Agent 1 eats up item o_{k+s}. Denote $l_{2,\text{II}}$ and $l_{3,\text{II}}$ the time span that Agent 1 eats item o_{k+s} and items $o_{j+1}, \cdots, o_k$, respectively.
- In Scenario III, the eating process is the same as Scenario II until $t_1 + l_{2,\text{II}}$. Denote the set of all agents by A. In proving Lemma 4, we divide A into four groups. That is, $A = \{\text{Agent } 1\} \cup A_2 \cup A_3 \cup A_4$, where A_2 is the set of agents that eat some of item o_{k+s} in Scenario I, A_3 is the set of agents that do not eat item o_{k+s} in Scenario I but would eat it in Scenario IV, and $A_4 = A \setminus (\{\text{Agent } 1\} \cup A_2 \cup A_3)$. For each agent $i \in A_2 \cup A_3$, denote the time period it eats item o_{k+s} after the moment t_1 in Scenario IV by $[t_1 + S_i, t_1 + E_i]$. In Scenarios III, we pause agent $i \in A_2 \cup A_3$ from eating in time $[t_1 + l_{2,\text{II}} + S_i, t_1 + l_{2,\text{II}} + E_i]$. Denote $l_{3,\text{III}}$ the time span that Agent 1 eats items $o_{j+1}, \cdots, o_k$.
- Scenario IV is a normal eating process when Agent 1 changes its preference order by placing item o_{k+s} after all items in $\overline{O}$. Denote $l_{3,\text{IV}}$ the time span that Agent 1 eats items $o_{j+1}, \cdots, o_k$ in Scenario IV. See Fig. 5.

Lemma 4. *Agent 1's utility in Scenario IV is no less than that in Scenario I.*

Proof. Clearly, Agent 1's utility received in the time period $[0, t_1]$ is the same in both Scenarios I and IV. Since its utility on eating item o_{k+s} is negligible, it suffices to compare its utility received by eating items $o_{j+1}, \ldots, o_k$. That is, to compare the value of $l_{3,\text{I}}$ and $l_{3,\text{IV}}$. First, due to Lemma 1, the volume of items $o_{j+1}, \ldots, o_k$ in Scenario II are no less than that in

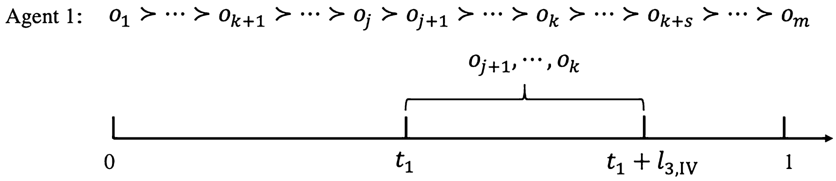

Fig. 5. Agent 1's preference order and the eating process in Scenario 4. Given agents' reports in which Agent 1 places item o_{k+s} after all items in $\overline{O}$, it is the normal scenario without any pause operation.

Scenario I at any time after the moment t_1. As such, it takes longer for the agents to eat up these items in Scenario II. Therefore, $l_{3,\mathrm{I}} \le l_{3,\mathrm{II}}$. Second, in view of Scenario II after the moment $t_1 + l_{2,\mathrm{II}}$, Scenario III is a pause scenario in which agents $i \in A_2 \cup A_3$ are paused from eating in time $[t_1 + l_{2,\mathrm{II}} + S_i, t_1 + l_{2,\mathrm{II}} + E_i]$. Hence, by Lemma 1, $l_{3,\mathrm{II}} \le l_{3,\mathrm{III}}$. Last, we compare the eating process in Scenario III after the moment $t_1 + l_{2,\mathrm{II}}$ and the eating process in Scenario IV after the moment t_1. They are the same except that agents $i \in A_2 \cup A_3$ are paused from eating in time $[t_1 + l_{2,\mathrm{II}} + S_i, t_1 + l_{2,\mathrm{II}} + E_i]$ in Scenario III and are eating item o_{k+s} in time $[t_1 + S_i, t_1 + E_i]$ in Scenario IV. Effectively, the agents' consumption of items $O \setminus \{o_{k+s}\}$ is the same in two scenarios. Therefore, $l_{3,\mathrm{III}} = l_{3,\mathrm{IV}}$. In conclusion, $l_{3,\mathrm{I}} \le l_{3,\mathrm{IV}}$. □

Now, we use Lemma 4 to prove Lemma 3.

Proof of Lemma 3. Given Lemma 2, if Agent 1 reports a preference order in which some items $o_{k+1}, \cdots, \emptyset_{k+s}$ in $\underline{O}$ are placed in front of items in $\overline{O}$, by repetitively moving the least preferred item o_{k+s} in $\underline{O}$ to the back of all items in $\overline{O}$, Agent 1's utility is non-decreasing. □

3.2. The incentive ratio upper bound when $n \ge m$

Denote u_1 and u'_1 Agent 1's utility in the truthful profile and the manipulation profile, respectively. Denote T and T' the moment by which all items in $\overline{O}$ are eaten up in the truthful profile and the manipulation profile, respectively. Denote $\tilde{T}$ and $\tilde{T}'$ the moment by which all items in $\overline{O}$ are eaten up while Agent 1 is paused all the time (or, say, Agent 1 is eliminated from the eating process) in the truthful profile and the manipulation profile, respectively.

When $n \ge m$, we prove the following theorem according to the value of T. That is, $0 < T < \frac{1}{2}$, $\frac{1}{2} \le T < \frac{2}{3}$, and $\frac{2}{3} \le T \le 1$.

Theorem 1. *When the number of agents is no less than that of items, the incentive ratio is upper bounded by $\frac{3}{2}$.*

When we ignore the ϵ terms for maintaining a strict ordering preference, obviously, by Lemma 3, $u_1 = T \le T' = u'_1$. By Lemma 2, $T' \le \tilde{T}'$. We also note that $\tilde{T} = \tilde{T}'$, as Agent 1's reports have no impact on the eating process when Agent 1 is eliminated. Therefore, we can prove our main result $u'_1 \le \frac{3}{2} u_1$ ($u'_1 \le 2u_1$) by showing that $\tilde{T} \le \frac{3}{2} T$ ($\tilde{T} \le 2T$). This way, we circumvent the requisition of figuring out Agent 1's best response strategies and only need to focus on how much more time it takes other agents to consume the items without the presence of Agent 1. This approach is used in proving Lemma 5, Case 1 of Lemma 6 and Case 1 of Theorem 2.

Lemma 5. *When $0 < T < \frac{1}{2}$, the Incentive Ratio is upper bounded by $\frac{3}{2}$.*

Proof. We will employ Lemma 2 and we are interested in a subset of $\overline{O}$. Denote $\overline{O}^* \subseteq \overline{O}$ the set of items that Agent 1 gets a positive fraction in the truthful profile. W.l.o.g., assume $\overline{O}^* = \{o^*_1, \cdots, o^*_{k^*}\}$. Denote l_j the length of time that Agent 1 spent on eating item o^*_j, $j = 1, \cdots, k^*$. That being said, item o^*_j is eaten up at the moment $\sum_{h=1}^{j} l_h$. Therefore, we have $T = \sum_{j=1}^{k^*} l_j$.

Since $0 < T < \frac{1}{2}$, and it takes at least time $\frac{1}{2}$ for two agents to eat up an item, we know that at the moment $\sum_{h=1}^{j} l_h$, there are at least three agents eating item o^*_j, $j = 1, \cdots, k^*$. So, apart from Agent 1, there are at least two other agents eating each of these items o^*_j at the moment they are eaten up.

Now let us consider the following process in which Agent 1 is eliminated. From the start to l_1, all agents eat items according to their preferences. At time l_1, pause all other agents but the agents who are eating item o^*_1. Since Agent 1 is eliminated, it will take them some extra time δ_1 to eat up item o^*_1. Agent 1 is absent for l_1 time so far and there are at least two of these agents, so $\delta_1 \le \frac{l_1}{2}$. At the moment $l_1 + \delta_1$, resume all agents' eating procedure, and pause all agents but the agents who are eating item o^*_2 at the moment $l_1 + \delta_1 + l_2$. For the same reason, it will take these agents an extra $\delta_2 \le \frac{l_2}{2}$ time to eat up the fraction due to Agent 1's absence. Repeating this process until all items in $\overline{O}^*$ are eaten up; at this moment all items in $\overline{O}/\overline{O}^*$ are eaten up as well. Therefore, we have that $\tilde{T} \le \sum_{j=1}^{k^*}(l_j + \delta_j) \le \sum_{j=1}^{k^*}(l_j + \frac{l_j}{2}) \le \frac{3}{2} T$. □

The second case is the most challenging one. It requires a fine-grained analysis.

Lemma 6. *When $\frac{1}{2} \le T < \frac{2}{3}$, the Incentive Ratio is upper bounded by 3/2.*

In this case, for every item $o^*_j \in \overline{O}^*$, if at least two other agents are eating the item with Agent 1 at the moment $\sum_{h=1}^{j} l_h$, then the lemma can be proved in the same way as Lemma 5. Therefore, assume that there exists an item that only one other agent is eating it with Agent 1 at the moment it is eaten up. Note that there could only exist one such item, as it takes at least time $\frac{1}{2}$ for two agents to eat up one item. Denote this item by $o^*_{k'}$ and the other agent by Agent 2. For ease of notation, let $t_1 = l_1 + \cdots + l_{k'-1}$, $t_2 = l_{k'}$, and $t_3 = l_{k'+1} + \cdots + l_{k^*}$. Then $t_1 + t_2 + t_3 = T$. To prove Lemma 6, we first prove the following auxiliary results.

Lemma 7. $t_1 + t_2 \ge \frac{1}{2}$, $t_3 < \frac{1}{6}$, $t_1 < \frac{1}{3}$, *and* $t_2 > 2t_3$.

Proof. Since only agents 1 and 2 are eating item $o^*_{k'}$, it must take at least time $\frac{1}{2}$ for them to eat it up. So $t_1 + t_2 \ge \frac{1}{2}$. In addition, $T < \frac{2}{3}$, so $t_3 = T - (t_1 + t_2) < \frac{1}{6}$. On the one hand, Agent 1 has been eating item $o^*_{k'}$ for time t_2, and Agent 2, even if it start eating item $o^*_{k'}$ from the beginning, has been eating it for time at most $t_1 + t_2$, we know that $t_2 + (t_1 + t_2) \ge 1$; on the other hand, $t_1 + t_2 \le T < \frac{2}{3}$, we conclude that $t_2 > \frac{1}{3}$ and $t_1 < \frac{1}{3}$. Therefore, we have $t_2 > \frac{1}{3} \ge 2t_3$. □

Denote O_1 the set of items that are eaten up on or before time t_1 in the truthful profile, O_2 the set of items that are eaten up in time interval $(t_1, t_1 + t_2]$ in the truthful profile, and O_3 the set of items that are eaten up in time interval $(t_1 + t_2, T]$ in the truthful profile. Note that these three sets contain all items that Agent 1 is interested in, i.e., $\overline{O} \subseteq O_1 \cup O_2 \cup O_3$.

Since $t_1 < \frac{1}{3}$, in time interval $[0, t_1]$, there are at least three agents eating any item that Agent 1 was eating, so the analysis for this interval is similar to that in Lemma 5.

Corollary 1. *If we eliminate Agent 1, all items in O_1 would be eaten up within time $\frac{3}{2}t_1$.*

For the set O_3, we would not obtain the same $\frac{3}{2}$ bound straightforwardly, but are able to obtain a slightly looser bound, which will be used together with some other approaches for handling O_2 to obtain an overall $\frac{3}{2}$ bound. We first show the following lemma.

Lemma 8. *In the normal scenario, for each item in O_3, at the moment that it is finished, there are at least two agents other than agents 1 and 2 who are eating the item.*

Proof. For each item $o_j \in O_3$, we prove the lemma in three possible cases.

- At the moment $t_1 + t_2$, no agent is eating item o_j. In this case, because $t_3 < \frac{1}{6}$ and the item o_j is eaten up within the interval $[t_1 + t_2, t_1 + t_2 + t_3]$, there must be at least six agents eating the item. Amongst these six agents, even if two of them are agents 1 and 2, there are another four agents.
- At the moment $t_1 + t_2$, only one agent is eating item o_j. In this case, even if this agent eats item o_j from the start, there are at least $1 - (t_1 + t_2)$ fraction of this item remaining, and it will be eaten up before the moment $t_1 + t_2 + t_3$, by $\frac{1-(t_1+t_2)}{t_3} = 1 + \frac{1-t}{t_3} > 1 + \frac{1-2/3}{1/6} = 3$, we know that there are at least four agents eating the item. Two of them might be agents 1 and 2. Even though, there are at least another two agents.
- At the moment $t_1 + t_2$, at least two agents are eating item o_j. Since agents 1 and 2 are eating item $o^*_{k'}$, these must be two other agents. □

If two agents were absent for some time, it will take another two agents the same amount of time to eat up the amount of items left over due to their absence. Therefore, a direct consequence of the above lemma is the following bound.

Corollary 2. *After eliminating Agent 1 from eating items in O_1, if we eliminate agents 1 and 2 from the moment $\frac{3}{2}t_1 + t_2$, all items in O_3 will be eaten up by at most an extra t_3 time.*

We now turn to prove Lemma 6 by combining these intermediate results and an analysis on item $o^*_{k'}$ and the set O_3. According to the high-level idea of our proofs described before, we will eliminate Agent 1 from eating item $o^*_{k'}$. This will delay the moment that item $o^*_{k'}$ is eaten up. More importantly, it will lead to two possible consequences.

Proof of Lemma 6. Case 1: The process of eating item $o^*_{k'}$ is extended, so some agents who eat items in O_3 will not continue eating their next item in O_3 and will take the chance to eat item $o^*_{k'}$ before the moment $\frac{3}{2}t_1 + t_2 + 2t_3$.

In the truthful profile, denote $\{s_1, \cdots, s_j\} \subseteq O_3$ the set of items that either Agent 1 or Agent 2 gets a positive fraction. For $h = 1, \cdots, j$, denote c_h and d_h the fraction of item s_h that agents 1 and 2 get, respectively.

In the manipulation profile (Agent 1 is eliminated, Agent 2 is eating item $o^*_{k'}$), denote z_h the moment at which item s_h is eaten up; denote Agent 3 the first agent except Agent 2 who comes to eat item $o^*_{k'}$ before the moment $\frac{3}{2}t_1 + t_2 + 2t_3$. In case there are multiple agents come to item $o^*_{k'}$ at the same time, pick one of them as Agent 3 randomly. Denote $\frac{3}{2}t_1 + t_2 + x$ the moment Agent 3 starts to eat item $o^*_{k'}$, where $0 \le x \le 2t_3$.[3]

Assume $1.5t_1 + t_2 + x \in (z_{w-1}, z_w]$, $w \le j$. At the moment $\frac{3}{2}t_1 + t_2 + x$, compare the truthful profile and the manipulation profile, due to the absence of Agent 1, item $o^*_{k'}$ is not eaten up in time. So both Agents 1 and 2 have not started eating items in O_3 at time $\frac{3}{2}t_1 + t_2$. So there is a delay of $\frac{1}{2}t_1 + \frac{\sum_{j=1}^{w-1} c_j + \sum_{j=1}^{w-1} d_j}{2}$. The $\frac{1}{2}$ in the second term is due to Lemma 8.

Now let Agents 2 and 3 eat item $o^*_{k'}$ and pause all other agents until the item is eaten up. This will take Agents 2 and 3 time $\frac{t_2 - x}{2}$, here t_2 is due to the absence of Agent 1, x is how much Agent 2 has eaten, the denominator 2 is due to Lemma 8.

Now at time $\frac{3t_1}{2} + t_2 + x + \frac{t_2 - x}{2}$, resume all agents except Agent 1. We calculate the amount of time after this moment when the items in O_3 will be eaten up in the manipulation profile. It consists of three categories. In the manipulation profile, when we pause the all other agents but the agents who are eating item in $\overline{O} \cap O_3$, this moment belongs to the first category. The time in the first category is at most $\frac{\sum_{h=w}^{j} d_h}{2}$. Note that $t_3 = \sum_{h=1}^{j} d_h$, the time in the first category is at most

$$\frac{t_3 - \sum_{h=1}^{w-1} d_h}{2}.$$

The second category is associated with Agent 2. At time $\frac{3t_1}{2} + t_2 + x + \frac{t_2 - x}{2}$, Agent 2 will begin to eat an item denoted by o_q. Then in the truthful profile, at the moment $t_1 + t_2 + x - \frac{\sum_{j=1}^{w-1} c_j + \sum_{j=1}^{w-1} d_j}{2}$ Agent 2 is eating item o_q. If $o_q \notin O_3$, it implies o_q will not be eaten up at moment $t_1 + t_2 + t_3$ in the truthful profile. Agent 2 will not affect the moment items in O_3 been eaten up in the manipulation profile anymore. The second category would be empty. If $o_q \in O_3$, in the manipulation profile, we would pause all other agents but the agents who are eating item o_q. This pause moment belongs to the second category. In the manipulation profile, before the moment $1.5t_1 + t_2 + x$, the total time Agent 2 has been absent for eating item in O_3 is $x - \frac{\sum_{h=1}^{w-1}(c_h + d_h)}{2}$. These items are $o_1, \ldots, o_{w_1}$ and o_q. The length of time Agent 2 has been absent for eating item o_q is $x - \frac{3\sum_{h=1}^{w-1} c_h + \sum_{h=1}^{w-1} d_h}{2}$. In the pause period, it will take agents who are eating o_q at most

$$\frac{x}{2} - \frac{3\sum_{h=1}^{w-1} c_h + \sum_{h=1}^{w-1} d_h}{4}$$

time, which is the upper bound of the time in the second category.

The other moments before O_3 have been eaten up belongs to third category. In this category, manipulation profile is identical to the truthful profile except Agent 1 is not eating. The time in this category would be

$$t_3 - \Big(x - \frac{\sum_{j=1}^{w-1} c_j + \sum_{j=1}^{w-1} d_j}{2}\Big).$$

Considering all three categories, we upper bound the time when all items in $O_1 \cup O_2 \cup O_3$ are eaten up as follows

$$\begin{aligned}
\tilde{T} \le & \Big(\frac{3t_1}{2} + t_2 + x + \frac{t_2 - x}{2}\Big) + \Big(\frac{t_3 - \sum_{h=1}^{w-1} d_h}{2}\Big) \\
& + \Big(\frac{x}{2} - \frac{3\sum_{h=1}^{w-1} c_h}{4} - \frac{\sum_{h=1}^{w-1} d_h}{4}\Big) \\
& + \Big(t_3 - x + \frac{c_1 + \ldots + c_{w-1} + d_1 + \ldots + d_{w-1}}{2}\Big) \\
= & \ 1.5t_1 + 1.5t_2 + 1.5t_3 - \frac{\sum_{h=1}^{w-1}(c_h + d_h)}{4} \\
\le & \ 1.5(t_1 + t_2 + t_3)
\end{aligned}$$

Case 2: Before time $1.5t_1 + t_2 + 2t_3$, the agents who eat items in O_3 will continue eating their next items in O_3 as they are more favorable than $o^*_{k'}$. In this case, we are not going to use the high-level idea presented before to show $\tilde{T} \le \frac{3}{2}T$.

[3] Note that only Agents 1 and 2 are eating item $o^*_{k'}$ in the normal scenario, so before the moment $\frac{3}{2}t_1 + t_2$, only Agent 2 is eating item $o^*_{k'}$ in the pause scenario. Hence, $x \ge 0$.

Instead, we will characterize an upper bound of Agent 1's utility when it uses the best response strategy. This upper bound will be partitioned into two quantities; one of which will be bounded using Lemma 2.

By using the best response strategy,

$$u_1' = u_1'|_{\leq \frac{3}{2}t_1+t_2+2t_3} + u_1'|_{> \frac{3}{2}t_1+t_2+2t_3},$$

where $u_1'|_{\leq \frac{3}{2}t_1+t_2+2t_3}$ and $u_1'|_{> \frac{3}{2}t_1+t_2+2t_3}$ denote Agent 1's utility gained before and after the moment $\frac{3}{2}t_1 + t_2 + 2t_3$, respectively. Obviously, $u_1'|_{\leq \frac{3}{2}t_1+t_2+2t_3} \leq \frac{3}{2}t_1 + t_2 + 2t_3$, where the equality holds if Agent 1 eats items in $\overline{O}$ from the start to time $\frac{3}{2}t_1 + t_2 + 2t_3$. Next, we upper bound $u_1'|_{> \frac{3}{2}t_1+t_2+2t_3}$.

In fact, the moment $\frac{3}{2}t_1 + t_2 + 2t_3$ is calibrated as it is the time needed to eat up items in O_1 and O_3, plus Agent 2 has been eating item $o_{k'}^*$ alone for t_2 time, when Agent 1 is eliminated. Now let us check out the items that would have been left over in the manipulation profile at the moment $\frac{3}{2}t_1 + t_2 + 2t_3$. By Lemma 2, we notice that the only item in $\overline{O}$ would possibly be still not eaten up yet, is $o_{k'}^*$. Due to the absence of Agent 1, it has t_2 left, but Agent 2 has been eaten it for an additional $2t_3$ time, so there would be $t_2 - 2t_3$ of item $o_{k'}^*$ left.

If $o_{k'}^*$ is not eaten up yet, Agent 2 must be eating it at the moment $1.5t_1 + t_2 + 2t_3$. To gain more utility, the optimal strategy for Agent 1 is to eat item $o_{k'}^*$ after this moment. Agent 1 can get at most $\frac{t_2-2t_3}{2}$ fraction for the existence of Agent 2. Therefore,

$$\begin{aligned} u_1' &= u_1'|_{\leq \frac{3}{2}t_1+t_2+2t_3} + u_1'|_{> \frac{3}{2}t_1+t_2+2t_3} \\ &\leq \frac{3}{2}t_1 + t_2 + 2t_3 + \frac{t_2 - 2t_3}{2} \\ &= \frac{3}{2}t_1 + \frac{3}{2}t_2 + t_3 \leq \frac{3}{2}T. \quad \square \end{aligned}$$

At last, we have the third case.

Lemma 9. *When $\frac{2}{3} \leq T \leq 1$, the Incentive Ratio is upper bounded by $\frac{3}{2}$.*

This is trivial, since the optimal utility u_1' is upper bound by $\frac{m}{n} \leq 1$. Hence the utility achieved in truthful profile is quite large compared to u_1'. Formally, $r^{PS} = \max \frac{u_1'}{u_1} \leq \frac{1}{\frac{2}{3}} \leq \frac{3}{2}$.

Combining Lemmas 5, 6, 9, we complete the proof of Theorem 1.

3.3. *The incentive ratio upper bound when $n < m$*

Theorem 2. *When the number of agents is less than that of items, the incentive ratio is upper bounded by 2.*

Proof. Recall that we denote $\overline{O}^* \subseteq \overline{O}$ the set of items that Agent 1 gets a positive fraction in the truthful profile. W.l.o.g., assume $\overline{O}^* = \{o_1^*, \dots, o_{k^*}^*\}$. Denote l_j the length of time that Agent 1 spends in eating item o_j^*, $j = 1, 2, 3, \dots, k^*$. Obviously, we have $u = \Sigma_{j=1}^{k^*} l_j$. According to the high-level idea to upper bound the highest utility Agent 1 can get, we divide all instances into two categories to analyze.

Case 1: Agent 1 does not monopolize any item in $\overline{O}^*$, i.e. at the moment $\Sigma_{j=1}^{h} l_j$, there are at least two agents eating item o_h^*. So apart from Agent 1, at least one other agent is eating each of these items in $\overline{O}^*$ at the moment they are eaten up. Now consider the process in which Agent 1 is eliminated from eating items. At time l_1, pause all agents that are not eating item o_1^*. Since Agent 1 is eliminated, it will take some extra time σ_1 to eat up o_1^*. As Agent 1 is absent for l_1 time and there is at least one other agent, $\sigma_1 \leq l_1$. Then at moment $l_1 + \sigma_1$, resume all agents' eating procedures, and pause all agents who are not eating item o_2^* at moment $l_1 + \sigma_1 + l_2$, then we can have o_2^* will be eaten up by the moment $l_1 + \sigma_1 + l_2 + \sigma_2$ and $\sigma_2 \leq l_2$, the reason is the same as aforementioned one. Repeat the process until all agents in $\overline{O}^*$ are eaten up, items in $\overline{O}/\overline{O}^*$ will be eaten up as well at this moment. Therefore, we have $\hat{T} \leq \Sigma_{j=1}^{k^*}(l_j + \sigma_j) \leq \Sigma_{j=1}^{k^*}(l_j + l_j) \leq 2u$.

Case 2: Agent 1 monopolizes some items in $\overline{O}^*$. W.l.o.g., we assume the set $\overline{O}_m^*$ of r ($r \leq k^*$) items is monopolized, i.e. for item $j \in \overline{O}_m^*$, $l_j = 1$. Denote O_{-1} the set of items that the other agents eat during the time period $[0, T]$ in the truthful profile. Through the same analysis of **Case 1**, for any item $j \in \overline{O}^*/\overline{O}_m^*$, it will take the other agents some extra time whose length is at most l_j apart from Agent 1. For items in $\overline{O}_m^*$, eliminating Agent 1 will lead to two possible consequences. The first is that some agents who eat items in O_{-1} will take the chance to eat items in $\overline{O}_m^*$ rather than continue eating the next item in O_{-1}. In this case, denote Agent 2 the first agent who comes to eat $o_m^* \in \overline{O}_m^*$ before all items in O_{-1} are eaten up.

If we pause all agents when Agent 2 is eating o_m^* and resume their eating procedure when o_m^* is eaten up. An extra time 1 will be incurred. W.l.o.g., assume a set $\overline{O}_f^* \subset \overline{O}_m^*$ of c $(c \le r)$ items in $\overline{O}_m^*$ will be eaten up before all items in O_{-1} are eaten up. If we repeat the aforementioned process for all of the c items, a total c period of time will be incurred for the same reason. Thus, we have that all items in $O_{-1} \bigcup \overline{O}^* / \overline{O}_m^* \bigcup \overline{O}_f^*$ will be eaten up before the moment $t' = T + \Sigma_{j \in \overline{O}_f^* \bigcup \overline{O}^* / \overline{O}_m^*} l_j$. All items in $\overline{O}_m^* / \overline{O}_f^*$ fall into the second category as they are less favorable than items in O_{-1} for the other agents. To conduct a further analysis on the impact of them, we characterize an upper bound of Agent 1's utility when it uses the best response strategy. Denote $u'|_{\le t'}$ and $u'|_{>t'}$ the utility Agent 1 gain before and after t', respectively. We have

$$u' = u'|_{\le t'} + u'|_{>t'}. \tag{1}$$

Obviously, $u'|_{\le t'} \le t'$. Since at moment t', only items in $\overline{O}_m^* / \overline{O}_f^*$ would possibly be items in $\overline{O}$ that are not eaten up yet. Thus, $u'|_{>t'} \le k - c = \Sigma_{j \in \overline{O}_m^* / \overline{O}_f^*} l_j$. Therefore,

$$u' = u'|_{\le t'} + u'|_{>t'} \le t' + \Sigma_{j \in \overline{O}_m^* / \overline{O}_f^*} l_j = 2 \cdot \Sigma_{j=1}^{k^*} l_j = 2u \quad \square \tag{2}$$

4. The tight bound examples

In this section, we provide two instances to show that the upper bounds obtained in Section 3 are tight.

4.1. When $n \ge m$

In the instance, the n agents' preferences are

$$\begin{aligned} 1:&\quad o_1 \succ o_2 \succ \cdots \succ o_{n-1} \succ o_n \\ 2:&\quad o_1 \succ o_2 \succ \cdots \succ o_{n-1} \succ o_n \\ i:&\quad o_2 \succ o_3 \succ \cdots \succ o_n \quad \succ o_1 \qquad i = 3, \cdots, n \end{aligned}$$

Agent 1 is interested in the first $\frac{n}{2} - 1$ items, i.e., $\overline{O} = \{o_1, \cdots, o_{\frac{n}{2}-1}\}$. Then in the truthful profile, $u_1 = \frac{1}{2}$. Agent 1 will only get half fraction of item 1.

By using the strategy $o_2 \succ o_3 \succ \cdots \succ o_{\frac{n}{2}-1} \succ o_1 \succ o_{\frac{n}{2}} \succ \cdots \succ o_n$. Agent 1 will get $\frac{1}{n-1}$ fraction of item 2 to $\frac{n}{2} - 1$ and $\frac{1}{4}$ fraction of item 1. Agent 1's utility becomes $u_1' = \frac{3}{4}$. So the ratio is $\frac{3}{2}$.

4.2. When $n < m$

In the instance, the n agents' preferences are:

$$\begin{aligned} 1:&\quad o_1 \succ o_2 \succ \ldots \succ o_{n-1} \succ o_n \succ \ldots \succ o_m \\ i:&\quad o_i \succ o_m \succ o_{m-1} \succ o_{m-2} \succ \ldots \succ o_{i+1} \succ o_{i-1} \succ \ldots \succ o_1 \qquad i = 2, \cdots, n \end{aligned}$$

Agent 1 is interested in the first n items, i.e., $\overline{O} = \{o_1, \ldots, o_n\}$. Then in the truthful profile, $u_1 = 1$.

By using the strategy $o_2 \succ o_3 \succ \ldots \succ o_n \succ o_1 \succ o_{n+1} \succ \ldots \succ o_m$. If $m \gg n$, Agent 1 will get $\frac{1}{2^{i-1}}$ fraction of item i. Then Agent 1's utility becomes $u' = \Sigma_{j=1}^{n} \frac{1}{2^{i-1}}$. When $n \to \infty$, u' approaches 2. So the ratio is 2.

5. Simulation results

In this section, we present numerical simulation on the extent to which an agent can increase its utility by unilateral manipulation in the Probabilistic Serial mechanism. While we investigated the theoretical incentive bounds of the mechanism in the worst-case framework, this section complements worst-case results by characterizing an agent's incentive to manipulate the PS mechanism in an average-case analysis sense.

5.1. Dichotomous preferences

Due to Lemma 1, we start with the case that Agent 1's preference values are dichotomous, i.e., either close to 1 or close to 0. Let the number of items that Agent 1 is interested be k. For each valid combination of $n \in [5, 10]$, $m \in [5, 10]$, and $k \in [3, 8]$, we generate 1,000 instances. In total, we generate $180,000$ instances. We construct an instance by generating uniformly, independently at random each Agent i's ordinal preferences, $i = 2, \cdots, n$. According to Lemma 3, Agent 1 will only consider reporting a preference that rearranges the items in $\overline{O}$. Therefore, for each instance, we enumerate Agent 1's all possible $k!$ preference orderings and record the largest ratio that its utility received by manipulation versus its utility received by truth-telling.

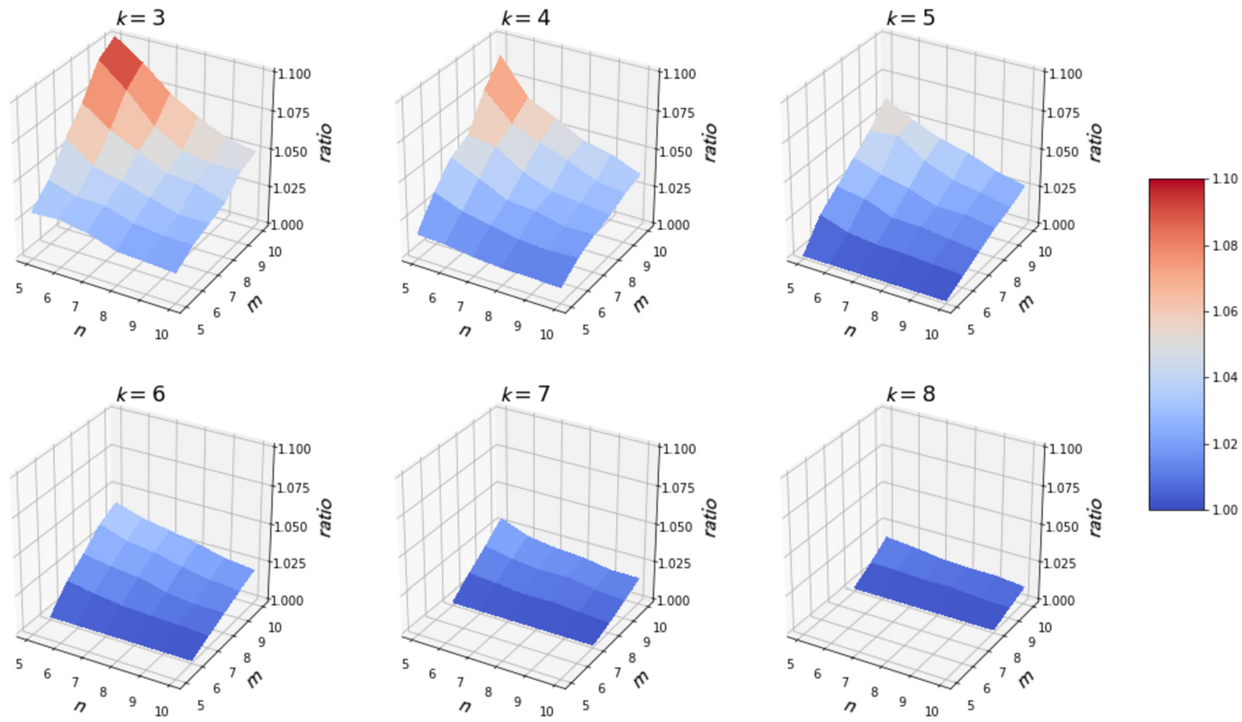

Fig. 6. Experimental evaluation. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Table 1
Fraction of instances whose ratio is greater than 1.05.

k	$\#\{(\succ_1, \succ_{-1}) \mid r_1 > 1.05\}$	$\#\ (\succ_1, \succ_{-1})$	%
$k = 3, n \in [5, 10], m \in [5, 10]$	10,709	36,000	29.75%
$k = 4, n \in [5, 10], m \in [5, 10]$	8,355	36,000	23.21%
$k = 5, n \in [5, 10], m \in [5, 10]$	5,489	36,000	15.25%
$k = 6, n \in [5, 10], m \in [6, 10]$	3,492	30,000	11.64%
$k = 7, n \in [5, 10], m \in [7, 10]$	1,790	24,000	7.46%
$k = 8, n \in [5, 10], m \in [8, 10]$	806	18,000	4.48%
Total	30,641	180,000	17.02%

Our simulation results are presented in Fig. 6. We can observe that the ratio is much smaller than the worst-case incentive ratio bounds of 1.5 and 2. Out of the 180,000 instances, 30,641 instances (17.02%) have a ratio larger than 1.05. A breakdown of these numbers is summarized in Table 1. In particular, for instances with a fixed n and m, the average ratio is decreasing in k. For instances with a fixed k, the average ratio is increasing in n or m.

5.2. Preference values follow distributions

In practice, an agent's preferences are not necessarily dichotomous. As such, we further relax Agent 1's values by randomly sampling them from probability distributions. For each combination of $n \in [4, 8]$ and $m \in [4, 8]$, we generate 1000 instances. That is $5 \times 5 \times 1,000 = 25,000$ instances. For each instance, we enumerate Agent 1's all possible $m!$ preference orderings and record the largest ratio that its utility received by manipulation versus its utility received by truth-telling. Therefore, in comparison to Fig. 6 in which the instances are still worst-case analysis to some extent by taking Lemma 1 and Theorem 2 into consideration, this experimental setting examines the ratio with more randomness. We perform experiments on three distributions, including Uniform distribution, Gaussian distribution, and Exponential distribution.

Fig. 7(a) shows the results when Agent 1's values, a_{1j}, are independently and identically drawn from the Uniform distribution U(0,1). As the figure shows, the ratios are much closer to 1 than the theoretical upper bounds of 1.5 and 2. In fact, only 1138 out of 25,000 ratios (4.55%) are larger than 1.05. Fig. 7(b) shows the results when a_{1j} are independently and identically drawn from the Gaussian distribution $N(0.5, 0.25)$, in which the values are truncated to be in $[0, 1]$. In this case, 772 out of 25,000 ratios (3.09%) are larger than 1.05. Fig. 7(c) shows the results when a_{1j} are independently and identically drawn from the Exponential distribution $Exp(0.2)$, in which the values are truncated to be in $[0, 1]$. In this case, 952 out of 25,000 ratios (3.81%) are larger than 1.05.

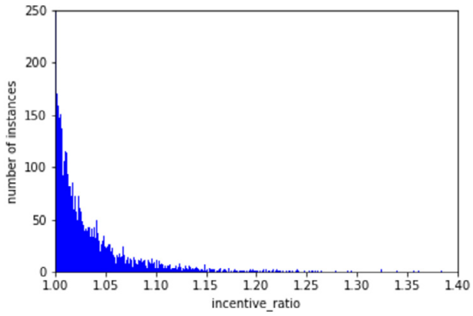

(a) Uniform distribution.

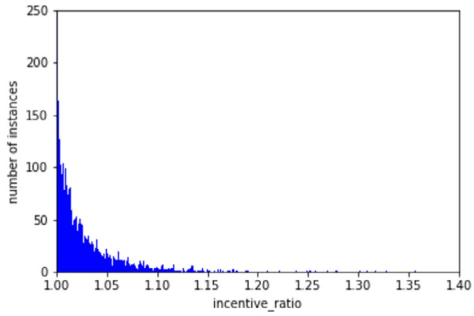

(b) Gaussian distribution.

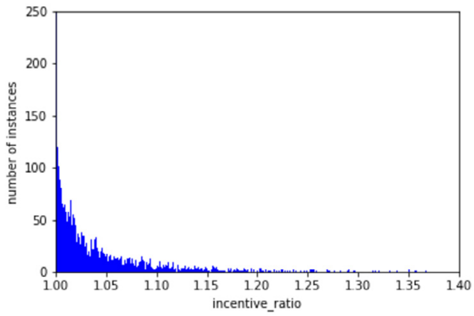

(c) Exponential distribution.

Fig. 7. Experiments.

6. Conclusion

Following the fact that the Probabilistic Serial mechanism is not incentive-compatible, in this paper, we examined the degree of an agent's incentive to manipulate the mechanism. In the form of incentive ratio, we showed that no agent is able to increase its utility by a large factor through strategic behaviors. This ratio is the strongest guarantee in the worst-case sense, by assuming that the agent is perfectly rational and has complete information about other agents' private information. In addition, we ran experiments to examine the manipulation incentive in an average-case sense [20]. The evaluation demonstrated that the utility-incremental ratio is much smaller than the theoretical worst-case bound. The evaluation bears out the supposition that Probabilistic Serial is approximately incentive-compatible in practice, even for small size instances. We hope that this work offered a better understanding of the robustness of the PS mechanism against manipulation, which is one step further than knowing that the mechanism is manipulable. An interesting next step is a formal analytical characterization of the average-case incentive ratio.

CRediT authorship contribution statement

Each of the authors has made an equal share of contributions in this paper.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgments

The authors thank the reviewers for their thoughtful comments and efforts towards improving the manuscript. Zihe Wang was supported by the Shanghai Sailing Program (Grant No. 18YF1407900) and the National NSFC (Grant No.

61806121). Jie Zhang was supported by a Leverhulme Trust Research Project Grant (2021–2024) and an EPSRC research grant (EP/W014912/1). Part of this work was done when Jie Zhang was visiting Peking University.

References

[1] A. Abdulkadiroglu, T. Sönmez, Random serial dictatorship and the core from random endowments in house allocation problems, Econometrica 66 (1998) 689–702.
[2] A. Abdulkadiroğlu, T. Sönmez, Matching markets: theory and practice, in: Advances in Economics and Econometrics (Tenth World Congress), 2013, pp. 3–47.
[3] Y. Azar, Y. Richter, The zero-one principle for switching networks, in: Proceedings of the 36th Annual ACM Symposium on Theory of Computing, 2004, pp. 64–71.
[4] H. Aziz, F. Brandt, M. Brill, The computational complexity of random serial dictatorship, Econ. Lett. 121 (2013) 341–345.
[5] H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, N. Narodytska, T. Walsh, Manipulating the probabilistic serial rule, in: Proceedings of International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, 2015, pp. 1451–1459.
[6] J.J. Bartholdi, J.B. Orlin, Single transferable vote resists strategic voting, Soc. Choice Welf. 8 (1991) 341–354.
[7] J.J. Bartholdi, C.A. Tovey, M.A. Trick, The computational difficulty of manipulating an election, Soc. Choice Welf. 6 (1989) 227–241.
[8] A. Bhalgat, D. Chakrabarty, S. Khanna, Social welfare in one-sided matching markets without money, in: APPROX-RANDOM, 2011, pp. 87–98.
[9] A. Bogomolnaia, E.J. Heo, Probabilistic assignment of objects: characterizing the serial rule, J. Econ. Theory 147 (2012) 2072–2082, https://doi.org/10.1016/j.jet.2012.05.013.
[10] A. Bogomolnaia, H. Moulin, A new solution to the random assignment problem, J. Econ. Theory 100 (2001) 295–328, https://doi.org/10.1006/jeth.2000.2710.
[11] A. Bogomolnaia, H. Moulin, Size versus fairness in the assignment problem, Games Econ. Behav. 90 (2015) 119–127, https://doi.org/10.1016/j.geb.2014.11.006.
[12] E. Budish, The combinatorial assignment problem: approximate competitive equilibrium from equal incomes, J. Polit. Econ. 119 (6) (2011) 1061–1103.
[13] Y.K. Che, F. Kojima, Asymptotic equivalence of probabilistic serial and random priority mechanisms, Econometrica 78 (2010) 1625–1672.
[14] N. Chen, X. Deng, B. Tang, H.R. Zhang, J. Zhang, Incentive ratio: a game theoretical analysis of market equilibria, Inf. Comput. 285 (2022) 104875, https://doi.org/10.1016/j.ic.2022.104875.
[15] N. Chen, X. Deng, H. Zhang, J. Zhang, Incentive ratios of Fisher markets, in: Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Proceedings, Part II, 2012, pp. 464–475.
[16] N. Chen, X. Deng, J. Zhang, How profitable are strategic behaviors in a market?, in: Algorithms - ESA 2011 - 19th Annual European Symposium, 2011, pp. 106–118.
[17] G. Christodoulou, A. Filos-Ratsikas, S.K.S. Frederiksen, P.W. Goldberg, J. Zhang, J. Zhang, Social welfare in one-sided matching mechanisms, in: N. Osman, C. Sierra (Eds.), Autonomous Agents and Multiagent Systems - AAMAS 2016, 2016, pp. 30–50.
[18] J. Davies, G. Katsirelos, N. Narodytska, T. Walsh, L. Xia, Complexity of and algorithms for the manipulation of Borda, Nanson's and Baldwin's voting rules, Artif. Intell. 217 (2014) 20–42, https://doi.org/10.1016/j.artint.2014.07.005.
[19] J. Davies, N. Narodytska, T. Walsh, Eliminating the weakest link: making manipulation intractable?, in: Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, 2012.
[20] X. Deng, Y. Gao, J. Zhang, Beyond the worst-case analysis of random priority: smoothed and average-case approximation ratios in mechanism design, Inf. Comput. 285 (2022) 104920, https://doi.org/10.1016/j.ic.2022.104920.
[21] S. Dughmi, A. Ghosh, Truthful assignment without money, in: ACM Conference on Electronic Commerce, 2010, pp. 325–334.
[22] Ö. Ekici, O. Kesten, An equilibrium analysis of the probabilistic serial mechanism, Int. J. Game Theory 45 (2016) 655–674, https://doi.org/10.1007/s00182-015-0475-9.
[23] Eric Budish, F.K.P.M. Yeon-Koo Che, Designing random allocation mechanisms: theory and applications, Am. Econ. Rev. 103 (2013) 585–623.
[24] A. Filos-Ratsikas, S.K.S. Frederiksen, J. Zhang, Social welfare in one-sided matchings: random priority and beyond, in: Algorithmic Game Theory, SAGT 2014, 2014, pp. 1–12.
[25] H. Hosseini, K. Larson, R. Cohen, Investigating the characteristics of one-sided matching mechanisms under various preferences and risk attitudes, Auton. Agents Multi-Agent Syst. 32 (2018) 534–567.
[26] D. Hugh-Jones, M. Kurino, C. Vanberg, An experimental study on the incentives of the probabilistic serial mechanism, Games Econ. Behav. 87 (2014) 367–380, https://doi.org/10.1016/j.geb.2014.06.001.
[27] A. Hylland, R. Zeckhauser, The efficient allocation of individuals to positions, J. Polit. Econ. 87 (1979) 293–314.
[28] Joseph Y. Halpern, R. Pass, L. Seeman, Decision theory with resource-bounded agents, Top. Cogn. Sci. 6 (2) (2014) 245–257.
[29] A.K. Katta, J. Sethuraman, A solution to the random assignment problem on the full preference domain, J. Econ. Theory 131 (2006) 231–250.
[30] O. Kesten, Why do popular mechanisms lack efficiency in random environments?, J. Econ. Theory 144 (2009) 2209–2226.
[31] T. Mennle, S. Seuken, An axiomatic approach to characterizing and relaxing strategyproofness of one-sided matching mechanisms, in: ACM-EC, 2014, pp. 37–38.
[32] T. Mennle, M. Weiss, B. Philipp, S. Seuken, The power of local manipulation strategies in assignment mechanisms, in: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, 2015, pp. 82–89.
[33] T. Sönmez, U. Ünver, Matching, allocation and exchange of discrete resources, in: Handbook of Social Economics 1A, 2011, pp. 781–852.
[34] L.G. Svensson, Strategy-proof allocation of indivisible goods, Soc. Choice Welf. 16 (1999) 557–567.
[35] L. Zhou, On a conjecture by Gale about one-sided matching problems, J. Econ. Theory 52 (1990) 123–135.