Bounding the Incentive Ratio of the Probabilistic Serial Rule

Bo Li
Hong Kong Polytechnic University
Hong Kong, China
comp-bo.li@polyu.edu.hk

Ankang Sun
Hong Kong Polytechnic University
Hong Kong, China
ankang.sun@polyu.edu.hk

Shiji Xing
Hong Kong Polytechnic University
Hong Kong, China
shi-ji.xing@connect.polyu.hk

ABSTRACT

Probabilistic Serial (PS) is a well-studied allocation rule used for distributing resources among multiple agents. Although it satisfies certain notable fairness and welfare properties, it is not truthful. This means that agents have incentives to misreport their preferences in order to influence the allocation in their favor. An interesting research question is to understand the extent to which an agent can gain from manipulation. A widely-accepted concept employed for this exploration is the *incentive ratio*, defined as the supreme ratio, across all instances of the problem, between the utility an agent obtains by employing an optimal manipulation strategy and the utility they receive when being truthful. Wang et al. [AAAI, 2020] examined the incentive ratio of PS for the setting when the number of items m equals the number of agents n and proved that the incentive ratio is 1.5. In this paper, we study the general scenario in which m and n can be arbitrary. We prove that in this case, the tight incentive ratio of PS is $2 - \frac{1}{2^{n-1}}$.

KEYWORDS

Probabilistic Serial; Manipulation; Incentive Ratio

ACM Reference Format:
Bo Li, Ankang Sun, and Shiji Xing. 2024. Bounding the Incentive Ratio of the Probabilistic Serial Rule. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

1 INTRODUCTION

The problem of fairly and efficiently distributing resources (called items in this paper) among agents has attracted significant attention in the field of computer science, artificial intelligence, and multi-agent systems [8, 10, 20, 21]. Many variants of this problem have been studied, and one of most popular settings is that the agents' preferences can be expressed as cardinal additive valuation functions over the items. An allocation rule decides who gets what based on the preferences reported by the agents. The fairness of an allocation is usually measured by, for example, envy-freeness [26] and proportionality [22], and the efficiency is measured by Pareto optimality [30]. An allocation is envy-free if every agent prefers her own bundle than any other agent's, and is proportional if every agent's utility is at least the average of her utility over all items. A Pareto optimal allocation ensures that nobody can get better off without hurting the other agents.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

Various allocation rules have been proposed and investigated towards meeting the desired properties. Several well-known ones include Picking Sequence [11], Envy-graph Procedure [20], Top Trading Cycles[1], and Probabilistic Serial [9]. In particular, Probabilistic Serial (SP) Rule is among the most prominent ones with appreciable properties on fairness and efficiency, including all the aforementioned envy-freeness, proportionality and Pareto optimality. The PS Rule was initially developed for the scenario where the number of agents and items are equal. Nonetheless, it can be naturally extended to the setting where there are more items than the agent [4, 18], preserving the fairness and efficiency guarantees.

The (extended) PS Rule operates as follows. First, each agent submits a ranking over all items. Then at the start of the PS Rule, every agent simultaneously begins to consume or “eat” her favorite item based on the reported ranking, all at the same rate. Multiple agents can eat one single item at the same time. Items are assumed to have a common capacity (without loss of generality, normalized to 1), and when an item is depleted, agents eating the item will move on their next available most preferred item. The PS Rule ends when all the items are exhausted. The bundle assigned to an agent is the set of fractional items that she has consumed in the procedure. In the context of allocating indivisible items, the fraction of an item consumed by some agent can also be interpreted as the probability of assigning the item to the agent.

Let us implement the PS Rule on an instance with two agents and three items denoted as $\{o_1, o_2, o_3\}$. The cardinal valuation of agents on every item is presented in Table 1 below where $\epsilon > 0$ is arbitrarily small.

	o_1	o_2	o_3
Agent 1	1	$1-\epsilon$	0
Agent 2	0	1	$1-\epsilon$

Table 1: Cardinal Valuations of Agents

The valuations induce the following ranking: agent 1 likes o_1 the most and prefers o_2 to o_3; agent 2 likes o_2 the most and prefers o_3 to o_1. Suppose that the consumption rate of both agents is 1, that is, it takes 1 unit of time for an agent to consume a whole item. We now simulate the procedure of the PS Rule. At the beginning, each agent starts to consume their most preferred item, that is, agent 1 starts to consume the item o_1 and agent 2 starts consuming o_2. At timestamp 1, both agents finish consuming their item; therefore, they start to consume their next most preferred item. For agent 1, her next most preferred item is o_2. However, o_2 has already been consumed, so agent 1 consumes the next available most preferred item, namely o_3. For agent 2, the next most preferred item for her is o_3, so she also starts consuming o_3. Now, both agents are consuming the same item, so after 0.5 unit of time o_3 is consumed. That is, at timestamp 1.5, all items are consumed and the mechanism ends. Upon termination, agent 1 (resp., agent 2) receives the entire o_1 (resp., o_2) and both agents receive 0.5 fraction of o_3.

On top of fairness and efficiency, *truthfulness* (also known as strategyproofness and incentive compatibility) is also highly desired when allocating resources. Unfortunately, the PS Rule is not truthful. An allocation rule is truthful if no agent can increase their utilities by lying. Thus, if an allocation rule is not truthful, strategic agents may have incentives to manipulate the resulting allocation by misreporting their preferences. This manipulation can compromise the overall performance of the mechanism.

Recall the example in Table 1. If both agents report their true preferences, we obtain the aforementioned allocation. Now, given that agent 2 continues to report truthfully, let us consider an alternative strategy of agent 1: she prefers o_2 the most and then prefers o_1 to o_3. Then the PS Rule returns the following allocation: agent 1 (resp., agent 2) receives the entire o_1 (resp., o_3) and both agents receive a fraction 0.5 of o_2. With this manipulation strategy, agent 1 is strictly better off. Note that agent 1's utility is still assessed by her true preference as shown in Table 1.

In general, retaining truthfulness together with other desired properties has been observed to be difficult or even impossible [2, 9]. For example, truthful, symmetric, and ex-ante Pareto optimality cannot be satisfied simultaneously as proved in [30], and truthful, equal treatment of equals, and ordinal efficient cannot be satisfied simultaneously as proved in [9]. Therefore, instead of pursing truthfulness on the top of fairness and efficiency, we want to understand the extent to which an agent can gain from manipulation in a particular allocation rule. On the one hand, such an understanding helps us analyze and predict agents' behaviors in different mechanisms, which can be further used to assess other quantitive properties such as price of anarchy [19]. On the other hand, if an agent cannot gain significantly from misreporting, it can be understood as the agent either finds it bothersome or lacks strong incentives to compute a manipulation strategy to manipulate the allocation. As a result, the mechanisms can still be considered satisfactory to some degree.

Therefore, in this paper, our goal is to quantify how much agents can benefit by misreporting in the PS Rule.

1.1 Our Contribution

To quantify the extent to which an agent can gain from manipulating in the PS Rule, we consider the commonly used notion, *incentive ratio*. The incentive ratio of an allocation rule is the supreme ratio between the utility an agent receives when employing the optimal manipulation strategy and the utility they receive when behaving truthfully, considering all possible allocation instances [14]. This problem has been partially addressed in the literature. Wang et al. [27] proved that the incentive ratio is 1.5 when the number of items m is at most the number of agents n. However, in many real-life resource allocation problems, the number of resources exceeds that of agents. Our work completely resolves this question by providing the tight bound of the incentive ratio of the PS Rule when the number of items and of agents can be arbitrary.

Our main contribution is the analysis of the tight incentive ratio of $2 - \frac{1}{2^{n-1}}$, where n represents the number of agents. We first observe that the worst-case (i.e., the agent can gain highest possible utility by manipulation) happens when the manipulator has dichotomous valuations; that is, the items can be categorized as large (with value 1) and small (with value 0). Then we present a reduction; that is, when bounding the incentive ratio, we can focus on the situation where all large items are ordered before small items in the reported ordinal preference of the manipulator. To upper bound the utility gain of the manipulator, we formalize a modified version of the PS Rule, where, except for the manipulator, other agents may stop consuming items before all items are exhausted. Based on the modified version of the PS Rule, we establish a sequence of upper bounds of the optimal value of the manipulator in the original PS Rule. The desired incentive ratio is established via computing the optimal objective value of a maximization problem.

1.2 Related works

Probabilistic serial rule. The Probabilistic Serial Rule was originally designed to address the house assignment problem [29], which involves n agents and n items, with each agent being assigned exactly one item. Bogomolnaia and Moulin [9] discusses the possible generalization of the PS Rule, where there can be more items than agents, but the constraint of each agent receiving only one item still applies. In a later work, Kojima [18] introduced the generalization where all items are consumed, which is the setting studied in our paper. The PS Rule can also be used to allocate indivisible items where the fraction of each item an agent receives is the probability she gets the item as a whole. This randomized algorithm ensures that the expected value of each agent's lottery is the best among all lotteries (bundles). However, the PS Rule has no ex-post fairness guarantee. Regarding this problem, Aziz [3] proposed a polynomial-time algorithm that has the same ex-ante behavior as the PS Rule while guaranteeing ex-post envy-freeness up to one item (EF1).

Although the PS Rule is susceptible to manipulation, computing the optimal manipulation strategy can be challenging. Aziz et al. [5] proved that computing the best response under lexicographic valuation can be done in polynomial time, but is NP-hard under general additive valuations. Aziz et al. [4] further proved that a pure Nash equilibrium always exists under the PS Rule, and verifying whether a given profile is a pure Nash equilibrium is coNP-complete.

Notions related to strategic behaviors. The concept of the incentive ratio was first used in the context of the Fisher Market [14] and then was generalized to other mechanisms, such as Resource Sharing [15] and Housing Markets [25]. Recently, Tao and Yang [24] and Xiao and Ling [28] respectively proved that the incentive ratio of the Round Robin algorithm for allocating indivisible resources is 2. Tao and Yang [24] also considered other mechanisms such as Maximum-Nash-Welfare and Envy-Graph Procedure. In particular, in this work, it is shown that the incentive ratio of MNW in general is unbounded. Later, Bei et al. [7] proved that if the items are homogeneously divisible, then MNW can achieve an incentive ratio of 2.

Several other notions related to strategic behaviors are considered in the literature, such as maximin strategy-proofness [12], risk-averse truthfulness [23], and price of anarchy [19]. Instead of considering individual manipulation potential, the price of anarchy concentrates on the potential loss of social welfare. Although it is generally challenging to achieve truthfulness, fairness, and efficiency simultaneously, there exist specific scenarios where favorable outcomes can be achieved. For example, regarding dichotomous valuations, Babaioff et al. [6] and Halpern et al. [16] respectively proposed truthful mechanisms that satisfy several fairness and social welfare properties.

Simultaneous work. It is worth mentioning that, in a recent independent work by Huang et al. [17], which is the extended journal version of the prior work by Wang et al. [27], the result of the incentive ratio of the PS Rule when $m = n$ is generalized to the setting when $m > n$. They proved a bound of 2, which is close to our result. Our paper involves a more fine-grained analysis, resulting in an additional improvement of $\frac{1}{2^{n-1}}$.

2 PRELIMINARIES

For any $t \in \mathbb{N}^+$, denote by $[t] = \{1, 2, \dots, t\}$. We consider the model of allocating a set $O = \{o_1, \dots, o_m\}$ of m divisible goods to a set $N = \{1, 2, \dots, n\}$ of n agents. Each agent i is associated with a strict *ordinal* preference $\succ_i^t$ over items, and $o_j \succ_i^t o_k$ refers to that agent i prefers o_j over o_k. Throughout the paper, $\succ_i^t$ always represents the true ordinal preference of agent i. Each agent i is also associated with additive cardinal valuation function $v_i : 2^O \to \mathbb{R}_{\geq 0}$, that is, $v_i(S) = \sum_{o \in S} v_i(o)$ for all $S \subseteq O$. For any $i \in [n]$, the cardinal valuation function $v_i(\cdot)$ is required to be consistent with $\succ_i^t$, in the sense that $v_i(o_t) \geq v_i(o_k)$ if and only if $o_t \succ_i^t o_k$. Moreover, for simplicity, we scale the cardinal valuations such that $\max_j v_i(o_j) = 1, \forall i \in [n]$. The underlying instance can then be represented as $\mathcal{I} = \langle N, O, \{\succ_i^t\}, \{v_i\} \rangle$.

An allocation is written in the form of a $n \times m$ matrix $\mathcal{X} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n)^T$ where the i-th row vector $\mathbf{X}_i = (x_{i1}, \dots, x_{im})$ refers to the (fraction of) items allocated to agent i, that is, a fraction x_{ij} of o_j is assigned to agent i in the allocation $\mathcal{X}$. With slight abuse of notation, $\mathbf{X}_i$ also refers to the set of fractional items assigned to agent i. Given an allocation $\mathcal{X}$, agent i's cardinal value on the received bundle $\mathbf{X}_i$ is $v_i(\mathbf{X}_i) = \sum_{j \in [m]} x_{ij} v_{ij}$. Since $\mathbf{X}_i$'s are vectors, operations such as summation and difference also apply to them. We will use the vector-wise calculation several times in the paper. We also let $|\mathbf{X}_i| = \sum_{j \in [m]} x_{ij}$, denoting the total size of the bundle $\mathbf{X}_i$.

In this work, following the convention, we study the extent to which individuals' value can be increased by *unilateral* deviation, in the sense that only one agent can misreport. Denote by $\text{PS}(\{\succ_j\}_{j \in [n]}, N, O)$ the Probabilistic Serial algorithm with agent set N, item set O and the reported ordinal preferences $\{\succ_j\}_{j \in [n]}$ (which may differ from their true preferences) as inputs. If the underlying inputs are clear from the context, we simply write PS. Since only one agent can manipulate, we often write the input of *PS* as $(\succ_i, \succ_{-i}, N, O)$ where $\succ_i$ is the reported preference of the agent i who can manipulate; $\succ_{-i}$ is the collection of reported preferences of all other agents (except agent i). Expression $\mathcal{X} = \text{PS}(\succ_i, \succ_{-i}, N, O)$ refers to that $\mathbf{X}$ is the allocation returned by PS with the corresponding input.

2.1 Incentive ratio

In order to quantify the extent to which individuals can benefit themselves through strategic play, we adopt the notion of *incentive ratio* ($\mathcal{R}$) [13, 14]. The incentive ratio of an allocation mechanism is defined as the supreme ratio over all problem instances between the largest possible value achieved by an agent through behaving strategically and the value received by that agent when she reveals the true information. The formal definition is presented below.

Definition 1 (Incentive ratio). *The incentive ratio of agent i with respect to an assignment rule $\mathcal{M}$ is defined as*

$$\mathcal{R}_i^{\mathcal{M}} = \sup_{\mathcal{I} = \langle N, O, \{v_i\}, \{\succ_i^t\} \rangle} \sup_{\succ_{-i}} \frac{\sup_{\succ_i} v_i(\mathcal{M}_i(\succ_i, \succ_{-i}), N, O)}{v_i(\mathcal{M}_i(\succ_i^t, \succ_{-i}), N, O)},$$

where $\mathcal{M}_i(\cdot)$ refers to the bundle assigned to agent i in the allocation returned by $\mathcal{M}$. The numerator is the largest possible value achieved by agent i when she unilaterally misreports the preference, while the denominator is the value of agent i when she reports truthfully. The incentive ratio of $\mathcal{M}$ is then $\mathcal{R}^{\mathcal{M}} = \max_{i \in [n]} \mathcal{R}_i^{\mathcal{M}}$.

We remark that for a given $i \in [n]$, the incentive ratio $\mathcal{R}_i^{\mathcal{M}}$ indeed does not depend on v_j and $\succ_j^t$ for all $j \neq i$, but for completeness, we keep these cardinal and ordinal valuations in the above fraction. By definition, the incentive ratio of an assignment rule $\mathcal{M}$ is at least one, and if it is equal to one, then no agent can increase their value in $\mathcal{M}$ by misreporting. Throughout the paper, without loss of generality, let agent 1 be the agent to manipulate and her true preference be $o_1 \succ_1^t o_2 \succ_1^t \cdots \succ_1^t o_m$. Moreover, as $\succ_{-1}^t$ does not affect $\mathcal{R}_i^{\mathcal{M}}$, thereafter, when bounding the incentive ratio, we further assume that agents except agent 1, report their true ordinal preferences.

We below use a concrete example to further illustrate the meaning of the incentive ratio. Let us again look the example presented in Introduction Section with two agents and three items $\{o_1, o_2, o_3\}$ and valuation functions are presented in Table 1. Note that agent 1 is the manipulator. Recall that if agent 1 reports her true preference, the returned allocation of PS is: $\mathbf{X}_1 = (1, 0, 0.5), \mathbf{X}_2 = (0, 1, 0.5)$. Then, the value of agent 1 is equal to $1 \times 1 + 0.5 \times 0 = 1$. Next we analyze the case where agent 1 manipulates her preference as $\succ_1'$: $o_2 \succ_1 o_1 \succ_1 o_3$. We demonstrate the running process of the PS Rule on this new preference profile $(\succ_1', \succ_2^t)$: At time 0, the most preferred item of both agents is o_2, and thus it takes 0.5 unit of time to finish consuming it. Then, the next preferred item of agent 1 is o_1, while that of agent 2 is o_3. Therefore, they will start to consume different items and finish them after 1 unit of time. The returned allocation is $\mathbf{X}_1' = (1, 0.5, 0), \mathbf{X}_2' = (0, 0.5, 1)$. This time, agent 1 receives a bundle of a value $1 \times 1 + 0.5 \times (1 - \epsilon) = 1.5 - 0.5\epsilon$. Indeed, this is the largest possible value agent 1 can get regardless of the reported preference of agent 2. Therefore, for this simple instance, the incentive ratio of agent 1 is $\frac{1.5 - 0.5\epsilon}{1} \approx 1.5$ when ϵ approaches 0. That is, agent 1 can increase her value by nearly 50% by misreporting her preference.

3 REDUCTIONS

In this section, we introduce two reductions that are helpful in finding the incentive ratio. The first one is an instance reduction that shrinks the space of worst-case instance, and hence enables us to concentrate on a specific class of instances. This reduction is established in [27], and states that when computing the incentive ratio of PS, it suffices to concentrate solely on *dichotomous preference* instance, that is, for any $o \in O$, the value of o of agent 1 $v_1(o)$ is either *close to* 1 or *close to* 0, and formally, for any $o \in O$, $v_1(o) \in \{1 - o(\epsilon), o(\epsilon)\}$ where $\epsilon > 0$ is arbitrarily small.

Lemma 1 (Lemma 1 of Wang et al. [27]). *Given any truthful profile $(\succ_1^t, \succ_{-1})$ and agent 1's cardinal valuation v_1 that is compatible with the ordering $\succ_1^t$, denote by ratio $c = \frac{u_1'}{u_1}$ where u_1' (resp. u_1) is agent 1's maximum value attainable by manipulation (resp. truthful reporting). Then one can always construct a corresponding dichotomous valuation b_i that is compatible with $\succ_1^t$, such that the ratio c is no less than before.*

Thereafter, agent 1's cardinal valuations is always assumed to be dichotomous. According to the dichotomous preference reduction, in agent 1's true ordinal preference, items that agent 1 is interested in (with a value close to 1) are ordered before the items that agent 1 is not interested in (with a value close to 0). Throughout the paper, denote by $\overline{O}$ the set of items in which agent 1 is interested, and assume $|\overline{O}| = k$; that is, $\overline{O} = \{o_1, o_2, \ldots, o_k\}$.

Now we present the second reduction, that is, when bounding the worst-case incentive ratio, we can focus on the case where the set of the first k items in agent 1's reported preference is identical to $\overline{O}$. We introduce an extra notation; given a $\succ_1$, let $C(\succ_1)$ be the number of items that are both in the first k elements in $\succ_1$ and in $\overline{O}$. The reduction replies on the following lemma.

Lemma 2. *For any $\succ_1, \succ_{-1}$, if $C(\succ_1) < k$, then there exists another ordinal preference $\succ_1'$ satisfying the following two properties:*

(1) $C(\succ_1') = k$;
(2) $v_1(\mathbf{X}_1) - v_1(\mathbf{X}_1') \leq O(\epsilon)$ where $\mathbf{X}_1$ and $\mathbf{X}_1'$ are the bundles assigned to agent 1 in $\mathsf{PS}(\succ_1, \succ_{-1}, N, O)$ and $\mathsf{PS}(\succ_1', \succ_{-1}, N, O)$ respectively.

Proof. First, we construct $\succ_1'$: we repeatedly find the first item in $\succ_1$ that is not in $\overline{O}$, and put this item to the last position on the preference profile, until the first k items in the profile all belong to $\overline{O}$. The resulting profile $\succ_1'$ ensures that $C(\succ_1') = k$, and the relative order of those k items remains the same as in $\succ_1$.

Before showing the second property, we will use a *paused allocation* to bridge the comparison between $\mathbf{X}_1$ and $\mathbf{X}_1'$. If an agent is *paused* amid an allocation, she stops consuming any items, while the other non-paused agents continue eating as usual. Once the agent's pause is over (or she is *resumed*), she will come back and select the highest-ranked item on her preference profile that is still available, following the PS Rule.

Now we construct the paused allocation, where agent 1 is the only agent that would be paused. Her bundle in this scenario is denoted as $\mathbf{X}_1^p$. The preference profile used by agent 1 in this scenario is $\succ_1'$, which contains $\overline{O}$ in its first k elements. To have an intuitive idea of the construction of the paused scenario, assume that we begin the paused allocation and the base allocation $\mathsf{PS}(\succ_1, \succ_{-1}, N, O)$ at the same time, and pause agent 1 whenever she is eating some item $o \notin \overline{O}$ in $\mathsf{PS}(\succ_1, \succ_{-1}, N, O)$. Formally, let t_1 be the earliest timestamp when agent 1 starts eating an item $o \in \overline{O}$ in the process of $\mathsf{PS}(\succ_1, \succ_{-1}, N, O)$, and let t_1' be the earliest time stamp after t_1 when agent 1 starts eating an item $o \notin \overline{O}$. Then, let t_2 be the first timestamp after t_1' when agent 1 starts eating an item $o \in \overline{O}$, etc. The sequence of timestamps will be used to describe the paused allocation: At the beginning of the allocation, agent 1 is paused. Then, at timestamp t_1, resume agent 1, and pause her again at timestamp t_1'. Repeat this pattern until the allocation ends or the items in $\overline{O}$ are exhausted. When there is no item left in $\overline{O}$, we pause agent 1 indefinitely. This regulation ensures that agent 1 only consumes valuable items in the (process of) paused allocation.

After the construction of the paused allocation, we will show that $v_1(\mathbf{X}_1)$ and $v_1(\mathbf{X}_1^p)$ differ by $O(\epsilon)$. First, we illustrate the following claim:

Claim 1. *At any timestamp, for any item, its remaining amount in the paused allocation is no less than that in the base allocation.*

Proof of Claim 1. Presume there exists a contradiction. Let t be the earliest time, when there exists an item o, whose remaining amount in the paused allocation is less than that in the base allocation. Since t is the earliest time, there must be more agents consuming o in the paused allocation within time $(t, t+\epsilon)$. Let the agent consuming o in the paused allocation but not consuming it in the base allocation be j. Suppose that agent j is consuming item o' in the base allocation at time t. We first discuss the situation where $j \neq 1$, then, either $o \succ_j o'$ or $o' \succ_j o$ since any agent other than agent 1 reports the same preference profile in the two allocations. Suppose it is the former case, then agent j not eating item o in the base allocation means that there is no item o left at time t in the base allocation. However, o is still available in the paused allocation, which contradicts the assumption that the remaining amount of o in the pause allocation at t is less. Taking into account the latter case, j not eating o' in the paused allocation means that o' is already exhausted at time t in the paused allocation but is still available in the base allocation, contradicting the assumption that t is the earliest time when some item has less amount left in the paused allocation. Hereby the proof for $j \neq 1$ is complete, and we consider the situation where $j = 1$. Since agent 1 is consuming items in the paused allocation, it means that she is consuming a valuable item in the base allocation at this time, that is, $o' \in \overline{O}$.Because o is also in $\overline{O}$ due to the construction of the paused allocation, we can derive that the relative order of those two items in the two profiles is consistent. That is, either $o \succ_1 o'$ or $o' \succ_1 o$ for both paused and base allocations. From here, we can use the same arguments as in the case where $j \neq 1$ to prove this situation. □

By Claim 1, it is not hard to see that agent 1 is either consuming a valued item or is paused in the paused allocation, and she will not be paused as long as $\overline{O}$ is not exhausted in the base allocation. Therefore, we can say that agent 1 spends the same amount of time consuming items in $\overline{O}$ in both allocations. Since the valuation is dichotomous, value of items $\overline{O}$ for agent 1 only differs by at most $O(\epsilon)$. Consequently, the difference between $v_1(\mathbf{X}_1)$ and $v_1(\mathbf{X}_1^p)$ is at most $O(\epsilon)$.

Now, to prove Lemma 2, it suffices to show $v_1(\mathbf{X}_1^p) - v_1(\mathbf{X}_1') \leq O(\epsilon)$. For convenience, we name $\mathsf{PS}(\succ_1', \succ_{-1}, N, O)$ the reduced allocation. It is not hard to see that the paused allocation is the reduced allocation with some pauses. The intuition here is to show that pausing does not increase agents' value. We show this by first presenting a similar claim.

Claim 2. *In the paused allocation, given any timestamp t, let $P(t)$ be the total amount of time that agent 1 is paused in the time range $(0, t)$. Given any timestamp t and item o, the remaining amount of o at time $t - P(t)$ in the reduced allocation is no less than its remaining amount at time t in the paused allocation.*

Proof of Claim 2. Assume for contradiction, t is the earliest time, when there exists an item o whose remaining amount in the reduced allocation at time $t - P(t)$ is less than that in the paused allocation at time t. Since t is the earliest time, there must be more agents consuming o in the reduced allocation within time $(t - P(t), t + \epsilon - P(t + \epsilon))$. This implies that agent 1 is not paused in the paused allocation at time t, otherwise the time range will have a length of 0. Let j be the agent consuming o at time $t - P(t)$ in the reduced allocation but consuming item o' in the paused allocation at t. Suppose $o \succ_j o'$, then o is unavailable at time t in the paused allocation. However, o is still available at time $t - P(t)$ in the reduced allocation, leading to a contradiction. For the case of $o' \succ_j o$, item o' is already exhausted at time $t - P(t)$ in the reduced allocation, meaning that t is not the earliest timestamp with the assumed property, a contradiction. □

We now continue proving Lemma 2. Suppose that the last timestamp when agent 1 is eating an item (which is valuable) in the paused scenario is t_p, that is, agent 1 spends $t_p - P(t_p)$ time consuming valuable items in the paused allocation. By Claim 2, we know that agent 1 is consuming valuable items from time 0 to $t_p - P(t_p)$ in the reduced allocation. Since the rest of the reduced allocation will not decrease her valuation, we can conclude that $v_1(\mathbf{X}_1^p) - v_1(\mathbf{X}_1') \leq O(\epsilon)$. As we already have $v_1(\mathbf{X}_1) - v_1(\mathbf{X}_1^p) \leq O(\epsilon)$, combining these two inequalities gives the second property in the statement of Lemma 2. □

Recall that we have scaled the valuation function of agent 1 such that $\max_{j \in [m]} v_1(o_j) = 1$. Thus, if agent 1 reports the true preference, she receives a value of at least $\frac{1}{n} \gg O(\epsilon)$, making the term $O(\epsilon)$ not affect the *incentive ratio*. Therefore, when bounding the worst-case incentive ratio, we can further assume that the first k items in agent 1's reporting are always (a permutation of) $\overline{O}$.

4 THE INCENTIVE RATIO OF PS RULE

In this section, we establish the incentive ratio of the PS rule. Recall that Wang et al. [27] proved that when $n \geq m$, the incentive ratio of the PS Rule is bounded by 1.5. As we target the ratio of $2 - \frac{1}{2^{n-1}}$, greater than 1.5 for any $n \geq 2$, it suffices to focus solely on the case of $m > n$. We first remark that the incentive ratio is monotonically non-decreasing regarding the number of items, that is, the incentive ratio of instances with $k+1$ items should be at least as large as that of instances with k items for all $k > n$. The reason behind this is that if on the worse case instance with the k item, we add a new item o^{k+1} such that o^{k+1} is the least preferred item for any agent in the new instance, then it is not hard to verify that the incentive ratio of the new instance with $k+1$ items should be no less than that of the k-item instance. As a consequence, when bounding the incentive ratio, we can further assume $m \gg n$.

In order to establish the tight incentive ratio, we propose the algorithm PS′ that can be viewed as a generalization of the canonical PS Rule. The algorithm PS′ takes as input a set of agents D, a set of fractional items E, and *partial* ordinal preferences (POP) of agents over E and returns an allocation of (a subset of) E over agents D. The size of an item $o \in E$ can be smaller than 1 and the POP allows agents to have an incomplete preference list, that is, not every $o \in E$ is required to appear on the reported preference list. Given an item $o \in E$, let $s(o)$ be the size of item o (contained in E) and naturally $s(o) \in [0, 1]$ for all $o \in E$.

The PS′ also utilizes the idea of the PS Rule, while different from the latter, if some item $o \in E$ does not appear in the reported POP list of agent i, then agent i never eats o. The formal description of PS′$(\succ_1, \succ_{-1}, D, E)$ is as follows (see Algorithm 1). We also use an example to demonstrate how PS′ executes.

Algorithm 1 PS′: The Modified Probabilistic Serial Rule

Input: Agent set D, item set E and POP profile $\{\succ_1, \succ_{-1}\}$.
1: All agents start to consume E at time 0 with the same rate of 1.
2: At any moment, each agent consumes the most preferred non-exhausted item appearing on her POP.
3: For an agent j, once all the items appearing in her POP are exhausted, she stops and receives the set of fractional items she consumed.
4: The algorithm ends if either all agents stop or all items are exhausted.

Example 1. *Consider an instance with two agents $D = \{1, 2\}$ and four items $E = \{o_1, o_2, o_3, o_4\}$ with sizes $s(o_1) = \frac{1}{2}$ and $s(o_j) = 1$ for $j \geq 2$. Suppose that the reported POP of agents 1 and 2 are $o_1 \succ_1 o_3$ and $o_1 \succ_2 o_3 \succ_2 o_4$, respectively. Then, at the end, agent 1 gets a fraction $\frac{1}{4}$ of o_1 and a fraction $\frac{1}{2}$ of o_3; agent 2 gets a fraction $\frac{1}{4}$ of o_1, a fraction $\frac{1}{2}$ of o_3 and the entire o_4. Item o_2 is left unassigned.*

We remark that at the termination of the canonical PS Rule, all items are exhausted or assigned, while in PS′, some item can be left unallocated.

4.1 The upper bound of the incentive ratio

In this subsection, we prove that incentive ratio of the PS rule is at most $2 - \frac{1}{2^{n-1}}$. We now fix an instance $\mathcal{I} = \langle N, O, \{\succ_i^t\}, \{v_i\} \rangle$ and a collection of other agents' reported preferences $\succ_{-1} = \{\succ_i\}_{i \geq 2}$. According to Lemma 1, we can focus on the case where function $v_1(\cdot)$ is dichotomous. Let T be the moment when $\overline{O}$ is eaten up in PS$(\succ_1^t, \succ_{-1}, N, O)$; recall that $\overline{O}$ is the set of items valued by agent 1. Moreover, denote by $\mathcal{X} = (\mathbf{X}_1, \ldots, \mathbf{X}_n)$ the partial allocation returned by PS$(\succ_1^t, \succ_{-1}, N, O)$ at moment T and by $\mathbf{X} = \sum_{i \in [n]} \mathbf{X}_i$ the set of items consumed in PS by moment T. Note that for vector $\mathbf{X} = (x_1, x_2, \ldots, x_m)$, every element x_j is at most one. For an item o_j, if $x_j = 1$, then we say o_j is an *integral* item of $\mathbf{X}$; otherwise, we say o_j is a *fractional* item of $\mathbf{X}$. It is worthwhile to note that in agent 1's true ordinal preference $\succ_1^t$, each item $o \in \overline{O}$ is ranked before any $o' \in O \setminus \overline{O}$. Accordingly, by the construction of $\mathcal{X}$, bundle $\mathbf{X}_1$ only contains items from $\overline{O}$. Moreover, as function v_1 is a dichotomous, we have $v_1(\mathbf{X}_1) = |\mathbf{X}_1| = T$, the received value of agent 1 when reporting the true preference. Note that since term $O(\epsilon)$ does not affect the incentive ratio, for simplicity, thereafter the dichotomous valuation function $v(\cdot)$ is assumed to be $v(o) \in \{0, 1\}$ for all o.

In the following, we establish upper bounds of T. Suppose that, for $\mathsf{PS}(\cdot, \succ_{-1}, N, O)$, agent 1's optimal strategy is to report $\succ_1^o$, and by doing so, she receives a value of $\mathsf{OPT}(1)$. Naturally, $\mathsf{OPT}(1) \geq T$. Our first step is to provide an upper bound of $\mathsf{OPT}(1)$, and the high-level idea is that if we can (virtually) let other agents stop at some moment earlier than T, then agent 1 would have opportunities to consume more of her valued items. Meanwhile, we also want to control or limit the additional value of agent 1 caused by the earlier stopping time of other agents. Therefore, we implement $\mathsf{PS}'(\cdot)$ in another instance $\mathcal{I}_1$ where the underlying items and agents are $\mathbf{X}$ and N, respectively. The cardinal valuation of agent 1 now becomes $v_1'(o) = 1$ for all o contained in $\mathbf{X}$, that is, instead of $\overline{O}$, agent 1 values every item in $\mathbf{X}$. For any $i \geq 2$, valuation v_i' remains unchanged, i.e., $v_i'(\cdot) = v_i(\cdot)$. The partial ordinal preferences of the agents $\{\succ_i^X\}_{i \in N}$ regarding $\mathcal{I}_1$ is constructed as follows:

- for agent 1, the partial ordinal preference $\succ_1^X$ contains all items of $\mathbf{X}$. Her favorite k items and their order are identical to that of $\succ_1^o$ and all other items of $\mathbf{X}$ are ordered arbitrarily afterward.
- for agent $i \geq 2$, let o_{j_i} be the item that agent i is eating at moment T in $\mathsf{PS}(\succ_1^t, \succ_{-1}, N, O)$. The partial ordinal preference $\succ_i^X$ only contains items $o_j \succ_i o_{j_i}$ and o_{j_i}, and their relative orders are consistent with $\succ_i$.

To help understand, we provide an example of $\succ_i^X$ for $i \geq 2$. Suppose $\mathbf{X}$ contains items $o_1, \ldots, o_5$ and agent i's preference over $\mathbf{X}$ is $o_5 \succ_i o_3 \succ_i o_1 \succ_i o_4 \succ_i o_2$. Item o_4 is the last one received by agent i, then $\succ_i^X$ becomes $o_5 \succ_i^X o_3 \succ_i^X o_1 \succ_i^X o_4$. Note that based on the definition of $\succ_i^X$, agent 2, in $\mathsf{PS}'(\cdot, \succ_{-1}^X, N, \mathbf{X})$, never consumes item o_2. In other words, when o_5, o_3, o_1, o_4 are exhausted, agent i stops. We now present propositions regarding the integral and fractional items of $\mathbf{X}$.

Proposition 1. *If item o_j appears in some $\succ_i^X$ and o_j is not the least preferred in $\succ_i^X$, then o_j is an integral item of* $\mathbf{X}$.

Proof. As o_j is not the least preferred in $\succ_i^X$, agent i must receive some fraction of another item $o_{j'}$ with $o_j \succ_i^X o_{j'}$ in $\mathbf{X}_i$. Note that item o_j must be completely consumed before agent i starts eating $o_{j'}$. Thus, $x_j = 1$ holds where x_j is the j-th element of vector $\mathbf{X}$, and in other words,o_j is an integral item of $\mathbf{X}$. □

Proposition 2. *Every item $o \in \overline{O}$ is an integral item of* $\mathbf{X}$.

Proof. The proof directly follows from Proposition 1 and the construction of $\overline{O}$ and of $\mathbf{X}$. □

Next, we analyze $\mathsf{PS}'(\succ_1^X, \succ_{-1}^X, N, \mathbf{X})$ and compare agent 1's value in the allocation returned by $\mathsf{PS}'(\succ_1^X, \succ_{-1}^X, N, \mathbf{X})$ and $\mathsf{OPT}(1)$.

Lemma 3. *In the allocation returned by* $\mathsf{PS}'(\succ_1^X, \succ_{-1}^X, N, \mathbf{X})$*, agent 1 has a value, with respect to v_1', of at least* $\mathsf{OPT}(1)$.

Proof. For ease of representation, in this proof, we use PS' and PS to represent $\mathsf{PS}'(\succ_1^X, \succ_{-1}^X, N, \mathbf{X})$ and $\mathsf{PS}(\succ_1^o, \succ_{-1}, N, O)$, respectively. According to Proposition 2, every $o \in \overline{O}$ is an integral element of $\mathbf{X}$, and agent 1 values every $o \in \mathbf{X}$. Then, it suffices to show that for any integral item o_j of $\mathbf{X}$, from time $t \geq 0$ to the moment where o_j is eaten (or exhausted), the remaining amount of o_j in PS' is at least the remaining amount in PS.

For the sake of a contraction, suppose t' is the earliest moment, after which there exists an integral item $o_{j'}$ whose remaining amount in PS' is less than its remaining amount in PS. Accordingly, for an arbitrarily small amount $\delta > 0$, at moment $t' + \delta$, the number of agents who are eating $o_{j'}$ in PS' is larger than that number in PS. Denote by i' the agent who, at moment $t' + \delta$, is eating $o_{j'}$ in PS' but not in PS. Then, suppose agent i', at moment $t' + \delta$, is eating another item $o_{j''}$ in PS. Recall that, at moment $t' + \delta$, the remaining amount of $o_{j'}$ is larger than zero, then, agent i' prefers $o_{j''}$ to $o_{j'}$ (i.e., $o_{j''} \succ_{i'} o_{j'}$); otherwise, agent i' should eat $o_{j'}$ at that moment. As for PS', agent i' is eating $o_{j'}$ at $t' + \delta$ in PS', which means, at that time, $o_{j''}$ is eaten up in PS'. However, $o_{j''}$ is not eaten up in PS at moment $t' + \delta$ as agent i' is eating that item. If we can prove $o_{j''}$ is also an integral item of $\mathbf{X}$, then t' should not be the earliest moment, providing a contradiction. As $o_{j'}$ appears in $\succ_{i'}^X$ and $o_{j''} \succ_{i'} o_{j'}$, item $o_{j''}$ is not the least preferred one in $\succ_{i'}^X$. Then, according to Proposition 1, $o_{j''}$ is also an integral item of $\mathbf{X}$, resulting in the desired contradiction. Up to here, we complete the proof. □

Denote by $\succ_1^{o(X)}$ the optimal strategy of agent 1 for $\mathsf{PS}'(\cdot, \succ_{-1}^X, N, \mathbf{X})$ and by $\mathsf{OPT}(2)$ the value of agent 1 when reporting $\succ_1^{o(X)}$ in $\mathsf{PS}'(\cdot, \succ_{-1}^X, N, \mathbf{X})$. It is not hard to see that $\mathsf{OPT}(2)$ is at least as large as the value of agent 1 in $\mathsf{PS}(\succ_1^X, \succ_{-1}^X, N, \mathbf{X})$, and based on Lemma 3, $\mathsf{OPT}(2) \geq \mathsf{OPT}(1)$ holds.

Proposition 3. $\mathsf{OPT}(2) \geq \mathsf{OPT}(1)$.

Proof. The proof follows directly from Lemma 3. □

We next present an upper bound of $\mathsf{OPT}(2)$ by implementing PS' to assign a subset of $\mathbf{X}$. Recall that $\mathbf{X} = (x_1, \ldots, x_m)$ and $\mathbf{X}_1 = (x_{11}, x_{12}, \ldots, x_{1m})$ represent both vectors and bundles. Our idea here is to decompose $\mathbf{X}$ into two parts: $\mathbf{X}_1$ and $\mathbf{X} - \mathbf{X}_1$ where $\mathbf{X} - \mathbf{X}_1$ refers to the vector (and hence the corresponding bundle) with its j-th element being $x_j - x_{ij}$ for all $j \in [m]$. Then we will directly assigns $\mathbf{X}_1$ to agent 1 for free and then use PS' to assign $\mathbf{X} - \mathbf{X}_1$, which can enhance agent 1's share on bundle $\mathbf{X} - \mathbf{X}_1$.

Formally, let $\mathcal{I}_2$ be an instance where the underlying items are $\mathbf{X} - \mathbf{X}_1$ and the underlying agents are N. The cardinal valuation function of each agent is identical to that in $\mathcal{I}_1$, and thus agent 1 values every $o \in \mathbf{X}$ (and hence every $o \in \mathbf{X} - \mathbf{X}_1$). Then, we present the following lemma.

Lemma 4. *Given any $\succ_1'$ over $\mathbf{X}$, for an integral item o of $\mathbf{X}$, from time $t \geq 0$ to the moment when o is eaten up, the remaining amount of o in $\mathsf{PS}'(\succ_1', \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$ is no more than the remaining amount in $\mathsf{PS}'(\succ_1', \succ_{-1}^X, N, \mathbf{X})$.*

Proof. If o is not included in $\mathbf{X} - \mathbf{X}_1$, then the statement is trivially true. Regarding the case where o is also included in $\mathbf{X} - \mathbf{X}_1$, the proof is similar to that of Lemma 3. □

Now we are ready to present an upper bound of $\mathsf{OPT}(2)$. Define $\mathsf{OPT}(3)$ be the maximum value that can be obtained by agent 1 in $\mathsf{PS}'(\cdot, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$.

Proposition 4. $\mathsf{OPT}(3) + T \geq \mathsf{OPT}(2)$.

Proof. Recall that $\succ_1^{o(X)}$ is the optimal strategy of agent 1 for $\mathsf{PS}'(\cdot, \succ_{-1}^X, N, \mathbf{X})$. For any $i \geq 2$, let l_i^* be the moment when agent i stops in $\mathsf{PS}'(\succ_1^{o(X)}, \succ_{-1}^X, N, \mathbf{X})$, and moreover, let p_i^* be the moment when agent i stops in $\mathsf{PS}'(\succ_1^{o(X)}, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$. According to Lemma 4, we have $l_i^* \geq p_i^*$ for all $i \geq 2$. Note that agent 1 values every item in $\mathbf{X}$, so $\mathbf{X}$ is eaten up in both $\mathsf{PS}'(\succ_1^{o(X)}, \succ_{-1}^X, N, \mathbf{X})$ and $\mathsf{PS}'(\succ_1^{o(X)}, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$. Consequently, the value of agent 1 is equal to the total size of $\mathbf{X}$, denoted as $|\mathbf{X}| = \sum_j x_{ij}$, minus the sum of the stopping time of agents $\{2, 3, \dots, n\}$. Thus, we have the following:

$$\begin{aligned}\mathsf{OPT}(2) &= |\mathbf{X}| - \sum_{i\geq 2} l_i^* \leq |\mathbf{X}| - \sum_{i \geq 2} p_i^* \\ &= |\mathbf{X}_1| + |\mathbf{X} - \mathbf{X}_1| - \sum_{i\geq 2} p_i^* \leq |\mathbf{X}_1| + \mathsf{OPT}(3) \\ &= v_1'(\mathbf{X}_1) + \mathsf{OPT}(3),\end{aligned}$$

where the last inequality transition is due to the fact that $|\mathbf{X} - \mathbf{X}_1| - \sum_{i \geq 2} p_i^*$ is the value of agent 1 when reporting $\succ_1^{o(X)}$ in $\mathsf{PS}'(\cdot, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$ and the last equality transition is due to the fact that in v_1', agent 1 values all items of $\mathbf{X}_1$. Moreover, as $T = v_1(\mathbf{X}_1) = |\mathbf{X}_1| = v_1'(\mathbf{X}_1)$, we then have $\mathsf{OPT}(3) + T \geq \mathsf{OPT}(2)$. □

After providing a sequence of upper bounds, at this stage, we find $\mathsf{OPT}(3) + T$, the target upper bound of $\mathsf{OPT}(1)$ and recall that $\mathsf{OPT}(1)$ is the value received by agent 1 in $\mathsf{PS}(\cdot, \succ_{-1}, N, O)$ when reporting the optimal strategy $\succ_1^o$. In what follows, we upper bound $\mathsf{OPT}(3)$, which then provides the desired upper bound ratio on the incentive ratio of the PS rule. Before presenting the main result, we first state a property of PS', that is, informally, at any moment t, eliminating an agent and the set of items currently assigned to her does not affect the assignment before t if one (virtually) re-runs the algorithm.

Lemma 5. *Given any ordinal preferences $\{\succ_i'\}$, a set of agents N, and a set of items E, suppose that at moment t, the partial assignment of $\mathsf{PS}'(\succ_1', \succ_{-1}', D, E)$, is $\mathbf{A}^t = (\mathbf{A}_1^t, \dots, \mathbf{A}_n^t)$. If we remove agent i and $\mathbf{A}_i^t$, then for $\mathsf{PS}'(\succ_{-i}, N \setminus \{i\}, E - \mathbf{A}_i^t)$, at moment t, the partial allocation is $\mathbf{B}^t = (\mathbf{B}_1^t, \dots, \mathbf{B}_{i-1}^t, \mathbf{B}_{i+1}^t, \dots, \mathbf{B}_n^t)$. It holds that $\mathbf{A}_j^t = \mathbf{B}_j^t$ for any $j \neq i$.*

Proof. For the sake of a contradiction, let t be the earliest moment after which $\mathbf{A}_{j^*}^t \neq \mathbf{B}_{j^*}^t$ for some $j^* \neq i$. Note that the agents $N \setminus \{i\}$ have identical ordinal preferences in these two scenarios. Accordingly, after the exhaustion of some item o, each agent $j \neq i$ will move to the same item in both scenarios. Therefore, if the allocation of an item o has changed at time t, it means that some agents have finished their previous item earlier, contradicting that t is the earliest moment that the difference appears. □

Now we are ready to present the main result of this paper: the incentive ratio of the PS rule is at most $2 - \frac{1}{2^{n-1}}$.

Theorem 1. *The incentive ratio of Probabilistic Serial is at most $2 - \frac{1}{2^{n-1}}$.*

Proof. We start from an arbitrary instance $\mathcal{I} = \langle N, O, \{\succ_i^t\}, \{v_i\}\rangle$ and an arbitrary collection of the reported preference of other agents $\succ_{-1}$, and show that $\mathsf{OPT}(3) + T \geq \mathsf{OPT}(1)$ where $\mathsf{OPT}(1)$ is the optimal value of agent 1 in $\mathsf{PS}(\cdot, \succ_{-1}, N, O)$. Note that the incentive ratio of PS is upper bounded by $\frac{\mathsf{OPT}(1)}{T}$ where T is the value of agent 1 when reporting truthfully in $\mathsf{PS}(\cdot, \succ_{-1}, N, O)$ and further, the incentive ratio of PS is at most $\frac{\mathsf{OPT}(3)}{T} + 1$. In what follows, we compute an upper bound of $\frac{\mathsf{OPT}(3)}{T}$ by formulating it as a maximization problem.

Let $\succ_1^{o^*}$ be the optimal strategy of agent 1 for $\mathsf{PS}'(\cdot, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$, and accordingly, agent 1 receives a value $\mathsf{OPT}(3)$ when reporting $\succ_1^{o^*}$ in $\mathsf{PS}'(\cdot, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$. Now we consider $\mathsf{PS}'(\succ_1^{o^*}, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$ and fix $S \subset N \setminus \{1\}$ as a set of agents. For each agent $i \in N$, let $\widetilde{t_i}$ be the moment when agent i stops. Define $p = \arg\max_{i\in S} \widetilde{t_i}$ and $\mathbf{A}_1^p$ the partial assignment of agent 1 at moment $\widetilde{t_p}$ in $\mathsf{PS}'(\succ_1^{o^*}, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$. By construction, we have the size $|\mathbf{A}_1^p| = \widetilde{t_p}$. Then, consider $\mathsf{PS}'(\succ_{-1}^X, N \setminus \{1\}, \mathbf{X} - \mathbf{X}_1 - \mathbf{A}_1^p)$ where agent 1 and her bundle $\mathbf{A}_1^p$ are removed. According to Lemma 5, for any agent $i \geq 2$ at any time $t \leq \widetilde{t_p}$, the assignment of $\mathsf{PS}'(\succ_1^{o^*}, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$ is identical to that of $\mathsf{PS}'(\succ_{-1}^X, N \setminus \{1\}, \mathbf{X} - \mathbf{X}_1 - \mathbf{A}_1^p)$. As a consequence, each agent $i \in S$, in $\mathsf{PS}'(\succ_{-1}^X, N \setminus \{1\}, \mathbf{X} - \mathbf{X}_1 - \mathbf{A}_1^p)$ also stops at moment $\widetilde{t_i}$, implying in $\mathsf{PS}'(\succ_{-1}^X, N \setminus \{1\}, \mathbf{X} - \mathbf{X}_1 - \mathbf{A}_1^p)$,

$$\sum_{i \in N \setminus \{S \cup \{1\}\}} \text{the stopping time of agent } i = |\mathbf{X} - \mathbf{X}_1 - \mathbf{A}_1^p| - \sum_{i\in S} \widetilde{t_i}.$$

Now consider $\mathsf{PS}'(\succ_{-1}^X, N \setminus \{1\}, \mathbf{X} - \mathbf{X}_1)$. By a similar argument to Lemma 4, the stopping time of each agent in $\mathsf{PS}'(\succ_{-1}^X, N \setminus \{1\}, \mathbf{X} - \mathbf{X}_1)$ should be no less than that in $\mathsf{PS}'(\succ_{-1}^X, N \setminus \{1\}, \mathbf{X} - \mathbf{X}_1 - \mathbf{A}_1^p)$ as the underlying items in the latter are a subset of the items in the former. Therefore, in $\mathsf{PS}'(\succ_{-1}^X, N \setminus \{1\}, \mathbf{X} - \mathbf{X}_1)$, the sum of stopping time over agents $N \setminus \{S \cup \{1\}\}$ is at least $|\mathbf{X} - \mathbf{X}_1 - \mathbf{A}_1^p| - \sum_{i\in S} \widetilde{t_i}$, and accordingly, the sum of the stopping time over agents in S is at most $\sum_{i \in S} \widetilde{t_i} + |\mathbf{A}_1^p|$ because the total size of the item is $|\mathbf{X} - \mathbf{X}_1|$ and agent 1 does not participate in $\mathsf{PS}'(\succ_{-1}^X, N \setminus \{1\}, \mathbf{X} - \mathbf{X}_1)$. By arguments similar to Lemma 5, each agent $i \in S$ should stop at moment T in $\mathsf{PS}'(\succ_{-1}^X, N \setminus \{1\}, \mathbf{X} - \mathbf{X}_1)$ due to the construction of $\mathbf{X}, \mathbf{X}_1$ and $\succ_{-1}^X$. Therefore, we have the following lower bound

$$\sum_{i\in S} \widetilde{t_i} + |\mathbf{A}_1^p| = \sum_{i\in S} \widetilde{t_i} + \widetilde{t_p} \geq |S|T, \tag{1}$$

which will act as constraints of a linear programming defined later on.

Recall that agent 1's optimal strategy is $\succ_1^{o^*}$ for $\mathsf{PS}'(\cdot, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$ and receives a value of $\mathsf{OPT}(3)$ when reporting $\succ_1^{o^*}$. By the construction described above, $\widetilde{t_i}$ is the stopping time of agent i in $\mathsf{PS}'(\succ_1^{o^*}, \succ_{-1}^X, N, \mathbf{X} - \mathbf{X}_1)$. As agent 1's partial ordinal preference list includes every item of $\mathbf{X} - \mathbf{X}_1$, then $\widetilde{t_1} \geq \widetilde{t_i}$ for all $i \in N$; note that $\mathsf{OPT}(3) = \widetilde{t_1}$. Without loss of generality, assume $\widetilde{t_2} \leq \widetilde{t_3} \leq \ldots \leq \widetilde{t_n} \leq \widetilde{t_1}$.

Next, we formulate a linear programming problem where the goal is to maximize the value of $\widetilde{t_1}$. The constraints of the problem are derived from inequalities (1); by taking $S = \{2\}$, we have $2\widetilde{t_2} \geq T$; by taking $S = \{2, 3\}$, we have $\widetilde{t_2} + 2\widetilde{t_3} \geq 2T$; by taking $S = \{2, 3, 4\}$, we have $\widetilde{t_2} + \widetilde{t_3} + 2\widetilde{t_4} \geq 3T$; and so on. Then, we have the following

LP formulation:

$$\begin{aligned}
\text{maximize} \quad & \widetilde{t}_1 \\
\text{subject to} \quad & 2\widetilde{t}_2 \geq T \\
& \widetilde{t}_2 + 2\widetilde{t}_3 \geq 2T \\
& \vdots \\
& \widetilde{t}_2 + \widetilde{t}_3 + \cdots + \widetilde{t}_{n-2} + 2\widetilde{t}_{n-1} \geq (n-2)T \\
& \widetilde{t}_2 + \widetilde{t}_3 + \cdots + \widetilde{t}_{n-1} + 2\widetilde{t}_n \leq (n-1)T \\
& \widetilde{t}_1 = \widetilde{t}_n \\
& \widetilde{t}_i \geq 0, \qquad \text{for } i \in N
\end{aligned}$$

where the third last constraint is due to the fact that the total size of $\mathbf{X} - \mathbf{X}_1$ is $(n-1)T$ and the second last constraint is due to the fact that $\mathbf{X} - \mathbf{X}_1$ can be eaten up by $N \setminus \{1\}$ under $\succ_{-1}^{X}$. By rearranging the first $n-2$ constraints, it is not hard to verify that the closed form of the optimal objective value is $(1 - \frac{1}{2^{n-1}})T$. Therefore, we have $\mathsf{OPT}(1) \leq \mathsf{OPT}(3) + T = \widetilde{t}_1 + T \leq (2 - \frac{1}{2^{n-1}})T$, which further implies the following,

$$\mathcal{R}^{\mathsf{PS}} \leq \frac{\mathsf{OPT}(1)}{T} \leq 2 - \frac{1}{2^{n-1}}.$$

Up to here, we complete the proof for the upper bound part. □

4.2 The matched lower bound example

In this subsection, we provide a concrete example, demonstrating that the upper bound derived in the previous section is indeed tight. We remark that our example is very similar to the lower bound instance presented in the simultaneous work [17]. For completeness, we also present our example and analyze the corresponding incentive ratio. Let us consider an instance with n agents and $2n$ items. The value of every item for agent 1 is specified in Table 2 where $\epsilon > 0$ is arbitrarily small.

o_1	o_2	$\cdots$	o_{n-1}	o_n
1	$1-\epsilon$	$\cdots$	$1-(n-2)\epsilon$	$1-(n-1)\epsilon$
o_{n+1}	o_{n+2}	$\cdots$	o_{2n-1}	o_{2n}
$(n-1)\epsilon$	$(n-2)\epsilon$	$\cdots$	ϵ	0

Table 2: The valuation function of agent 1

According to the valuation function, the set of items valued by agent 1 is $\overline{O} = \bigcup_{j=1}^{n}\{o_j\}$. Moreover, for any $i \neq 1$, the ordinal preference of agents i is defined as follows:

$$\succ_i\colon o_i \succ_i o_{2n} \succ_i o_{2n-1} \succ_i \cdots \succ_i o_1$$

The ordinal preference $\succ_i$, with $i \neq 1$, is roughly the inverse of $\succ_1^t$, and the only difference is that o_i is now ranked as the top choice.

Recall that each item has size 1 and agents have a consumption rate of 1. Then the PS Rule will terminate at time 2. Let $\mathbf{X}_1$ be the bundle allocated to agent 1 when she reports truthfully and let $\mathbf{X}_1'$ be the set of items assigned to her when she reports the optimal ordinal preference with respect to the PS Rule. First consider the case where agent 1 reports the truthful preference. From time 0 to 1, every agent i consumes o_i. Afterwards, the items left are not valued by agent 1. Therefore at the termination, agent 1's value $v_1(\mathbf{X}_1)$ is equal to $1 + O(\epsilon)$.

We now consider a manipulation strategy of agent 1, shown as follows,

$$\succ_1'\colon o_2 \succ_1 o_3 \succ_1 \cdots o_n \succ_1 o_1 \succ_1 o_{n+1} \succ_1 o_{n+2} \succ_1 \cdots$$

Given the reported preference profile $(\succ_1', \succ_{-1})$, agents 1 and 2 consume o_2 simultaneously from time 0 to $\frac{1}{2}$. After that, agent 1 moves on to her next preferred item, o_3. However, agent 3 has been consuming o_3 since the beginning and has already consumed half of it. Therefore, from time $\frac{1}{2}$ to $\frac{3}{4}$, agent 1 and agent 3 share o_3, with agent 1 consuming only a quarter of it. Agent 1 then moves on to consume o_4, which only has a quarter left, and this pattern continues. The consumption of o_n by agent 1 ends at time $\frac{1}{2} + \frac{1}{4} + \cdots + \frac{1}{2^{n-1}} = 1 - \frac{1}{2^{n-1}}$. After that, agent 1 starts consuming o_1 and we will prove that no other agents will compete with agent 1 on o_1, so that agent 1 consumes all of o_1. According to the ordinal preferences $\{\succ_i\}_{i\neq 1}$, every agent, except agent 1, must complete the consumption of the final n items before moving on to o_1. We now show that by time 2, the total size of the last n items consumed by agents $2, \ldots, n$ is $n - \frac{1}{2^{n-1}}$, less than n.

Note that agent 2 starts consuming the last n items at time $\frac{1}{2}$, and thus she consumes a size of $1 + \frac{1}{2}$ of the last n items by time 2. Similarly, the size of the last n items consumed (by time 2) by agent 3 is $1 + \frac{1}{4}$, and that of agent 4 is $1 + \frac{1}{8}$. It can be verified that the total size of the last n items consumed by agents $2, \ldots, n$ is equal to $(n-1) + 1 - \frac{1}{2^{n-1}} = n - \frac{1}{2^{n-1}}$, indicating that agent 1 can consume all o_1 before any other agents move on to o_1. Therefore, agent 1 can consumes the entire o_1 and receives a value of $v_1(\mathbf{X}_1') = 2 - \frac{1}{2^{n-1}} + O(\epsilon) \to 2 - \frac{1}{2^{n-1}}$ as $\epsilon \to 0$, which implies that the incentive ratio of the constructed instance is at least $2 - \frac{1}{2^{n-1}}$. This further indicates that the upper bound established in the previous section is tight.

5 CONCLUSION

In this paper, we proved that the incentive ratio of the PS Rule for the general setting where the number of items and agents can be arbitrary is equal to $2 - \frac{1}{2^{n-1}}$. Our results indicate that the incentive ratio of the PS Rule is only bounded by the number of agents and not by the number of items. For further directions, it would be intriguing to explore the possibility of transforming a PS problem into a round-robin allocation problem. Additionally, it would be interesting to discover another allocation rule with a smaller incentive ratio while maintaining attributes such as efficiency and fairness.

ACKNOWLEDGMENTS

The authors are ordered alphabetically. This work is funded by NSFC (No. 62102333), HKSAR RGC (No. PolyU 25211321) and GDSTC (No. 2023A1515010592).

REFERENCES

[1] Atila Abdulkadiroğlu and Tayfun Sönmez. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66, 3 (1998), 689–701.

[2] Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. 2017. Truthful Allocation Mechanisms Without Payments: Characterization and Implications on Fairness. In *EC 17*, Constantinos Daskalakis, Moshe Babaioff, and Hervé Moulin (Eds.). ACM, 545–562.

[3] Haris Aziz. 2020. Simultaneously Achieving Ex-ante and Ex-post Fairness. In *WINE 2020 (Lecture Notes in Computer Science, Vol. 12495)*, Xujin Chen, Nikolai Gravin, Martin Hoefer, and Ruta Mehta (Eds.). Springer, 341–355.

[4] Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Nina Narodytska, and Toby Walsh. 2015. Equilibria Under the Probabilistic Serial Rule. In *IJCAI 2015*, Qiang Yang and Michael J. Wooldridge (Eds.). AAAI Press, 1105–1112.

[5] Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Nina Narodytska, and Toby Walsh. 2015. Manipulating the Probabilistic Serial Rule. In *AAMAS 2015*, Gerhard Weiss, Pinar Yolum, Rafael H. Bordini, and Edith Elkind (Eds.). ACM, 1451–1459.

[6] Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2021. Fair and truthful mechanisms for dichotomous valuations. In *AAAI 2021*, Vol. 35. 5119–5126.

[7] Xiaohui Bei, Biaoshuai Tao, Jiajun Wu, and Mingwei Yang. 2023. The Incentive Guarantees Behind Nash Welfare in Divisible Resources Allocation. *CoRR* abs/2308.08903 (2023). arXiv:2308.08903

[8] Ivona Bezáková and Varsha Dani. 2005. Allocating indivisible goods. *SIGecom Exch.* 5, 3 (2005), 11–18.

[9] Anna Bogomolnaia and Hervé Moulin. 2001. A New Solution to the Random Assignment Problem. *J. Econ. Theory* 100, 2 (2001), 295–328.

[10] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. 2016. Fair Allocation of Indivisible Goods. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 284–310.

[11] Sylvain Bouveret and Jérôme Lang. 2014. Manipulating picking sequences.. In *ECAI*, Vol. 14. 141–146.

[12] Steven J Brams, Michael A Jones, Christian Klamler, et al. 2006. Better ways to cut a cake. *Not. Am. Math. Soc.* 53, 11 (2006), 1314–1321.

[13] Ning Chen, Xiaotie Deng, Hongyang Zhang, and Jie Zhang. 2012. Incentive Ratios of Fisher Markets. In *ICALP 2012 (Lecture Notes in Computer Science, Vol. 7392)*, Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts, and Roger Wattenhofer (Eds.). Springer, 464–475.

[14] Ning Chen, Xiaotie Deng, and Jie Zhang. 2011. How Profitable Are Strategic Behaviors in a Market?. In *ESA 2011 (Lecture Notes in Computer Science, Vol. 6942)*, Camil Demetrescu and Magnús M. Halldórsson (Eds.). Springer, 106–118.

[15] Yukun Cheng, Xiaotie Deng, and Yuhao Li. 2020. Tightening Up the Incentive Ratio for Resource Sharing Over the Rings. In *2020 IEEE*. IEEE, 127–136.

[16] Daniel Halpern, Ariel D Procaccia, Alexandros Psomas, and Nisarg Shah. 2020. Fair division with binary valuations: One rule to rule them all. In *WINE 2020*. Springer, 370–383.

[17] Haoqiang Huang, Zihe Wang, Zhide Wei, and Jie Zhang. 2024. Bounded incentives in manipulating the probabilistic serial rule. *J. Comput. Syst. Sci.* 140 (2024), 103491.

[18] Fuhito Kojima. 2009. Random assignment of multiple indivisible objects. *Math. Soc. Sci.* 57, 1 (2009), 134–142.

[19] Elias Koutsoupias and Christos H. Papadimitriou. 1999. Worst-case Equilibria. In *STACS (Lecture Notes in Computer Science, Vol. 1563)*. Springer, 404–413.

[20] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *EC 2004*, Jack S. Breese, Joan Feigenbaum, and Margo I. Seltzer (Eds.). ACM, 125–131.

[21] Benjamin Plaut and Tim Roughgarden. 2018. Almost Envy-Freeness with General Valuations. In *SODA 2018*, Artur Czumaj (Ed.). SIAM, 2584–2603.

[22] Hugo Steinhaus. 1948. The problem of fair division. *Econometrica* 16 (1948), 101–104.

[23] Biaoshuai Tao. 2022. On existence of truthful fair cake cutting mechanisms. In *EC 2022*. 404–434.

[24] Biaoshuai Tao and Mingwei Yang. 2023. Incentive Ratios for Fairly Allocating Indivisible Goods: Simple Mechanisms Prevail. *CoRR* abs/2306.15920 (2023). arXiv:2306.15920

[25] Taiki Todo. 2019. Analysis of Incentive Ratio in Top-Trading-Cycles Algorithms. In *JSAI 2019*. The Japanese Society for Artificial Intelligence, 2F1E304–2F1E304.

[26] Hal R Varian. 1974. Equity, envy, and efficiency. *J. Econ. Theory* 9, 1 (1974), 63–91.

[27] Zihe Wang, Zhide Wei, and Jie Zhang. 2020. Bounded Incentives in Manipulating the Probabilistic Serial Rule. In *AAAI 2020*. AAAI Press, 2276–2283.

[28] Mingyu Xiao and Jiaxing Ling. 2020. Algorithms for Manipulating Sequential Allocation. In *AAAI*. AAAI Press, 2302–2309.

[29] Lin Zhou. 1990. On a conjecture by Gale about one-sided matching problems. *J. Econ. Theory* 52, 1 (1990), 123–135.

[30] Lin Zhou. 1991. Inefficiency of strategy-proof allocation mechanisms in pure exchange economies. *Social Choice and Welfare* 8 (1991), 247–254.