

Social Welfare in One-Sided Matching Mechanisms ^{*}

George Christodoulou¹, Aris Filos-Ratsikas², Søren Kristoffer Stiil Frederiksen³, Paul W. Goldberg², Jie Zhang², and Jinshan Zhang¹

¹ Department of Computer Science, University of Liverpool, UK
`{gchristo,jinshan.zhang}@liv.ac.uk`

² Department of Computer Science, University of Oxford, UK
`{aris.filos-ratsikas,paul.goldberg,jie.zhang}@cs.ox.ac.uk`

³ Department of Computer Science, Aarhus University, Denmark
`ssf@cs.au.dk`

Abstract. We study the Price of Anarchy of mechanisms for the well-known problem of one-sided matching, or house allocation, with respect to the social welfare objective. We consider both ordinal mechanisms, where agents submit preference lists over the items, and cardinal mechanisms, where agents may submit numerical values for the items being allocated. We present a general lower bound of $\Omega(\sqrt{n})$ on the Price of Anarchy, which applies to *all* mechanisms. We show that two well-known mechanisms, Probabilistic Serial, and Random Priority, achieve a matching upper bound. We extend our lower bound to the Price of Stability of a large class of mechanisms that satisfy a common proportionality property, and show stronger bounds on the Price of Anarchy of all *deterministic* mechanisms.

1 Introduction

One-sided matching (also called the house allocation problem) is the fundamental problem of assigning items to agents, such that each agent receives exactly one item. It has numerous applications, such as assigning workers to shifts, students to courses or patients to doctor appointments. In this setting, agents are often asked to provide *ordinal preferences*, i.e. preference lists, or rankings of the items. We assume that underlying these ordinal preferences, agents have numerical values specifying how much they value each item [18]. In game-theoretic terms,

^{*} George Christodoulou, Paul W. Goldberg and Jinshan Zhang were supported by the EPSRC grant EP/K01000X/1. Paul W. Goldberg was supported by COST Action IC1205. Aris Filos-Ratsikas and Jie Zhang were supported by the ERC Advanced Grant 321171 (ALGAME). Aris Filos-Ratsikas and Søren K.S. Frederiksen acknowledge support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, within which part of this work was performed and support from the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council.

these are the agents' von Neumann-Morgenstern utility functions [27] and the associated preferences are often referred to as *cardinal preferences*.

A *mechanism* is a function that maps agents' valuations to matchings. However, agents are rational strategic entities that might not always report their valuations truthfully; they may misreport their values if that results in a better matching (from their own perspective). Assuming the agents report their valuations strategically to maximize their utilities, it is of interest to study the *Nash equilibria* of the induced game, i.e. strategy profiles from which no agent wishes to unilaterally deviate.

A natural objective for the designer is to choose the matching that maximizes the *social welfare*, i.e. the sum of agents' valuations for the items they are matched with, which is the most prominent measure of aggregate utility in the literature. Given the strategic nature of the agents, we are interested in mechanisms that maximize the social welfare *in the equilibrium*. We use the standard measure of equilibrium inefficiency, the *Price of Anarchy* [22], that compares the maximum social welfare attainable in any matching with the *worst-case* social welfare that can be achieved in an equilibrium.

We evaluate the efficiency of a mechanism with respect to the Price of Anarchy of the induced game. We study both deterministic and randomized mechanisms: in the latter case the output is a probability mixture over matchings, instead of a single matching. We are interested in the class of *cardinal* mechanisms, which use cardinal preferences, and generalize the ordinal mechanisms.

Note that our setting involves no monetary transfers and generally falls under the umbrella of *approximate mechanism design without money* [24]. In general settings without money, one has to fix a canonical representation of the valuations. A common approach in the literature is to consider the *unit-sum* normalization, i.e. each agent has a total value of 1 for all the items. We obtain results for unit-sum valuations, and extend most of these to another common normalization, *unit-range*.

1.1 Our results

In Section 3 we bound the inefficiency of the two best-known mechanisms in the matching literature, *Probabilistic Serial* and *Random Priority*. In particular, for n agents and n items, the Price of Anarchy is $O(\sqrt{n})$. In Section 4 we complement this with a *matching* lower bound (i.e. $\Omega(\sqrt{n})$) that applies to *all* cardinal (randomized) mechanisms. As a result, we conclude that these two *ordinal* mechanisms (ones that compute matchings that only depend on preference orderings) are optimal. These results suggest that it does not help a welfare maximizer to ask agents to report more than the ordinal preferences.

We separately consider *deterministic* mechanisms and in Section 4 prove that their Price of Anarchy is $\Omega(n^2)$, even for cardinal mechanisms. This shows that randomization is necessary for non-trivial worst-case efficiency guarantees.

In Section 5, we extend our results to more general solutions concepts as well as the case of incomplete information. Finally, in Section 6, we prove that under a mild "proportionality" property, our lower bound of $\Omega(\sqrt{n})$ extends to the

Price of Stability, a more optimistic measure of efficiency [3], which strengthens the negative results even further. Additionally, we discuss how our results extend to the other common normalization in the literature, *unit-range* [2, 15, 28].

1.2 Discussion and related work

The one-sided matching problem was introduced in [18] and has been studied extensively ever since (see [1] for a recent overview). Over the years, several different mechanisms have been proposed with various desirable properties related to truthfulness, fairness and economic efficiency with Probabilistic Serial [1, 7–9] and Random Priority [1, 2, 4, 9, 15, 23] being the two prominent examples.

As mentioned earlier, in settings without money, one needs to represent the valuations in some canonical way. A common approach is the *unit-sum* normalization, i.e. each agent has a total value of 1 for all the items. Intuitively, this normalization means that each agent has equal influence within the mechanism and her values can be interpreted as “scrip money” that she uses to acquire items. The unit-sum representation is standard for social welfare maximization in many settings without money including fair division, cake cutting and resource allocation [10, 11, 15, 16] among others. Moreover, without any normalization, non-trivial Price of Anarchy bounds cannot be achieved by any mechanism.

The objective of social welfare maximization for one-sided matching problems has been studied before in the literature, but mainly for truthful mechanisms [2, 15]. Our lower bounds are more general, since they apply to *all* mechanisms, not just truthful ones. In particular, our lower bound on the Price of Anarchy of all mechanisms generalizes the corresponding bound for truthful mechanisms in [15]. Note that Random Priority is truthful (truth-telling is a dominant strategy equilibrium) but it has other equilibria as well; we observe that the welfare guarantees of the mechanism hold for all equilibria, not just the truth-telling ones. Similar approaches have been made for truthful mechanisms like the second price auction in settings with money.

While given our general lower bound, proving a matching upper bound for Random Priority is enough to establish tightness, it is still important to know what the welfare guarantees of Probabilistic Serial are, given that it is arguably the most popular one-sided matching mechanism. The mechanism was introduced by [9] and since then, it has been in the center of attention of the matching literature, with related work on characterizations [17, 20], extensions [19], strategic aspects [21] and hardness of manipulation [6]. Somewhat surprisingly, the Nash equilibria of the mechanism were only recently studied. Aziz et al. [5] prove that the mechanism has pure Nash equilibria while Ekici and Kesten [14] study the *ordinal* equilibria of the mechanism and prove that the desirable properties of the mechanism are not necessarily satisfied for those profiles.

Another, somewhat different recent branch of study considers ordinal measures of efficiency instead of social welfare maximization, under the assumption that agents’ preferences are only expressed through preference orderings over items. Bhalgat et al. [8] study the approximation ratio of matching mechanisms, when the objective is maximization of *ordinal social welfare*, a notion of efficiency

that they define based solely on ordinal information. Other measures of efficiency for one-sided matchings were also studied in Krysta et al. [23], where the authors design truthful mechanisms to approximate the size of a maximum cardinality (or maximum agent weight) Pareto-optimal matching and in Chakrabarty and Swamy [12] where the authors consider the rank approximation as the measure of efficiency. While interesting, these measures of efficiency do not accurately encapsulate the socially desired outcome the way that social welfare does, especially since an underlying cardinal utility structure is part of the setting [9, 18, 27, 28]. Our results actually suggest that in order to achieve the optimal welfare guarantees, one does not even need to elicit this utility structure; agents can only be asked to report preference orderings, which is arguably more appealing.

Finally, we point out that our work is in a sense analogous to the literature that studies the Price of Anarchy in item-bidding auctions (e.g. see [13, 26] and references therein) for settings without money. Furthermore, the extension of our results to very general solution concepts (coarse correlated equilibria) and settings of incomplete information (Bayes-Nash equilibria) is somehow reminiscent of the *smoothness* framework [25] for games. While our results are not proven using the smoothness condition, our extension technique is similar in spirit.

2 Preliminaries

Let $N = \{1, \dots, n\}$ be a finite set of agents and $A = \{1, \dots, n\}$ be a finite set of indivisible items. An *allocation* is a matching of agents to items, that is, an assignment of items to agents where each agent gets assigned exactly one item. We can view an allocation μ as a permutation vector $(\mu_1, \mu_2, \dots, \mu_n)$ where μ_i is the unique item matched with agent i . Let O be the set of all allocations. Each agent i has a valuation function $u_i : A \rightarrow \mathbb{R}$ mapping items to real numbers. Valuation functions are considered to be well-defined modulo positive affine transformations, that is, for item $j : j \rightarrow \alpha u_i(j) + \beta$ is considered to be an alternative representation of the same valuation function u_i . Given this, we fix the canonical representation of u_i to be *unit-sum*, that is $\sum_j u_i(j) = 1$, with $u_i(j) \geq 0$ for all i, j . Equivalently, we can consider valuation functions as *valuation vectors* $u_i = (u_{i1}, u_{i2}, \dots, u_{in})$ and let V be the set of all valuation vectors of an agent. Let $\mathbf{u} = (u_1, u_2, \dots, u_n)$ denote a typical *valuation profile* and let V^n be the set of all valuation profiles with n agents.

We consider *strategic agents* who might have incentives to misreport their valuations. We define $\mathbf{s} = (s_1, s_2, \dots, s_n)$ to be a pure strategy profile, where s_i is the *reported* valuation vector of agent i . We will use \mathbf{s}_{-i} to denote the strategy profile without the i th coordinate and hence $\mathbf{s} = (s_i, \mathbf{s}_{-i})$ is an alternative way to denote a strategy profile. A *direct revelation mechanism* without money is a function $M : V^n \rightarrow O$ mapping *reported* valuation profiles to matchings. For a randomized mechanism, we define M to be a random map $M : V^n \rightarrow O$. Let $M_i(\mathbf{s})$ denote the restriction of the outcome of the mechanism to the i 'th coordinate, which is the item assigned to agent i by the mechanism. For randomized mechanisms, we let $p_{ij}^{M, \mathbf{s}} = \Pr[M_i(\mathbf{s}) = j]$ and $p_i^{M, \mathbf{s}} = (p_{i1}^{M, \mathbf{s}}, \dots, p_{in}^{M, \mathbf{s}})$. When it

is clear from the context, we drop one or both of the superscripts from the terms $p_{ij}^{M,\mathbf{s}}$. The utility of an agent from the outcome of a deterministic mechanism M on input strategy profile \mathbf{s} is simply $u_i(M_i(\mathbf{s}))$. For randomized mechanisms, an agent's utility is $\mathbb{E}[u_i(M_i(\mathbf{s}))] = \sum_{j=1}^n p_{ij}^{M,\mathbf{s}} u_{ij}$.

A subclass of mechanisms that are of particular interest is that of *ordinal mechanisms*. Informally, ordinal mechanisms operate solely based on the *ordering* of items induced by the valuation functions and not the actual numerical values themselves, while cardinal mechanisms take those numerical values into account. Formally, a mechanism M is *ordinal* if for any strategy profiles \mathbf{s}, \mathbf{s}' such that for all agents i and for all items j, ℓ , $s_{ij} < s_{i\ell} \Leftrightarrow s'_{ij} < s'_{i\ell}$, it holds that $M(\mathbf{s}) = M(\mathbf{s}')$. A mechanism for which the above does not necessarily hold is *cardinal*. Equivalently, the strategy space of ordinal mechanisms is the set of all permutations of n items instead of the space of valuation functions V^n . A strategy s_i of agent i is a *preference ordering* of items (a_1, a_2, \dots, a_n) where $a_\ell \succ a_k$ for $\ell < k$. We will write $j \succ_i j'$ to denote that agent i prefers item j to item j' according to her true valuation function and $j \succ_{s_i} j'$ to denote that she prefers item j to item j' according to her strategy s_i . When it is clear from the context, we abuse the notation slightly and let u_i denote the truth-telling strategy of agent i , even when the mechanism is ordinal. Note that agents can be indifferent between items and hence the preference order can be a weak ordering.

Two properties of interest are *anonymity* and *neutrality*. A mechanism is anonymous if the output is invariant under renamings of the agents and neutral if the output is invariant under relabeling of the objects.

An *equilibrium* is a strategy profile in which no agent has an incentive to deviate to a different strategy. First, we will focus on the concept of *pure Nash equilibrium*, formally

Definition 1. *A strategy profile \mathbf{s} is a pure Nash equilibrium if $u_i(M_i(\mathbf{s})) \geq u_i(M_i(s'_i, s_{-i}))$ for all agents i , and pure deviating strategies s'_i .*

In Section 5, we extend our results to more general equilibrium notions as well as the setting of incomplete information, where agents' values are drawn from known distributions.

Let $S_{\mathbf{u}}^M$ denote the set of all pure Nash equilibria of mechanism M under truthful valuation profile \mathbf{u} . The measure of efficiency that we will use is the *pure Price of Anarchy*,

$$PoA(M) = \sup_{\mathbf{u} \in V^n} \frac{SW_{OPT}(\mathbf{u})}{\min_{\mathbf{s} \in S_{\mathbf{u}}^M} SW_M(\mathbf{u}, \mathbf{s})}$$

where $SW_M(\mathbf{u}, \mathbf{s}) = \sum_{i=1}^n \mathbb{E}[u_i(M_i(\mathbf{s}))]$ is the expected *social welfare* of mechanism M on strategy profile \mathbf{s} under true valuation profile \mathbf{u} , and $SW_{OPT}(\mathbf{u}) = \max_{\mu \in \mathcal{O}} \sum_{i=1}^n u_i(\mu_i)$ is the social welfare of the optimal matching. Let $OPT(\mathbf{u})$ be the optimal matching on profile \mathbf{u} and let $OPT_i(\mathbf{u})$ be the restriction to the i th coordinate.

3 Price of Anarchy Guarantees

In this section, we prove the (pure) Price of Anarchy guarantees of Probabilistic Serial and Random Priority. Together with our lower bound in the next section, the results establish that both mechanisms are optimal among all mechanisms for the problem.

Probabilistic Serial

First, we consider *Probabilistic Serial*, which we abbreviate to *PS*. Informally, the mechanism is the following. Each item can be viewed as an infinitely divisible good that all agents can consume at unit speed during the unit time interval $[0, 1]$. Initially each agent consumes her most preferred item (or one of her most preferred items in case of ties) until the item is entirely consumed. Then, the agent moves on to consume the item on top of her preference list, among items that have not yet been entirely consumed. The mechanism terminates when all items have been entirely consumed. The fraction p_{ij} of item j consumed by agent i is then interpreted as the probability that agent i will be matched with item j under the mechanism.

We prove that the Price of Anarchy of *PS* is $O(\sqrt{n})$. Aziz et al. [5] proved that *PS* has pure Nash equilibria, so it makes sense to consider the pure Price of Anarchy; we will extend the result to the coarse correlated Price of Anarchy and the Bayesian Price of Anarchy in Section 5.

We start with the following two lemmas, which prove that in a pure Nash equilibrium of the mechanism an agent's utility cannot be much worse than what her utility would be if she were consuming the item she is matched with in the optimal allocation from the beginning of the mechanism until the item is entirely consumed. Let $t_j(\mathbf{s})$ be the time when item j is entirely consumed on profile \mathbf{s} under *PS*(\mathbf{s}).

Lemma 1. *Let \mathbf{s} be any strategy profile and let s_i^* be any strategy such that $j \succ_{s_i^*} \ell$ for all items $\ell \neq j$, i.e. agent i places item j on top of her preference list. Then it holds that $t_j(s_i^*, \mathbf{s}_{-i}) \geq \frac{1}{4} \cdot t_j(\mathbf{s})$.*

Proof. For ease of notation, let $\mathbf{s}^* = (s_i^*, \mathbf{s}_{-i})$. Obviously, if $j \succ_{s_i^*} \ell$ for all $\ell \neq j$ and since all other agents' reports are fixed, $t_j(\mathbf{s}^*) = t_j(\mathbf{s})$ and the statement of the lemma holds. Hence, we will assume that there exists some item $j' \neq j$ such that $j' \succ_{s_i^*} j$.

First, note that if agent i is the only one consuming item j for the duration of the mechanism, then $t_j(\mathbf{s}^*) = 1$ and we are done. Hence, assume that at least one other agent consumes item j at some point, and let τ be the time when the first agent besides agent i starts consuming item j in \mathbf{s}^* . Obviously, $t_j(\mathbf{s}^*) > \tau$, therefore if $\tau \geq \frac{1}{4} \cdot t_j(\mathbf{s})$ then $t_j(\mathbf{s}^*) \geq \frac{1}{4} \cdot t_j(\mathbf{s})$ and we are done. So assume that $\tau < \frac{1}{4} \cdot t_j(\mathbf{s})$. Next observe that in the interval $[\tau, t_j(\mathbf{s}^*)]$, agent i can consume at most half of what remains of item i because there exists at least one other agent consuming the item for the same duration. Overall, agent i 's consumption is at

most $\frac{1}{2} + \frac{1}{4}t_j(\mathbf{s})$ so at least $\frac{1}{2} - \frac{1}{4}t_j(\mathbf{s})$ of the item will be consumed by the rest of the agents.

Now consider all agents other than i in profile \mathbf{s} and let α be the amount of item j that they have consumed by time $t_j(\mathbf{s})$. Notice that the total consumption speed of an item is non-decreasing in time which means in particular that for any $0 \leq \beta \leq 1$, agents other than i need at least $\beta t_j(\mathbf{s})$ time to consume $\alpha \cdot \beta$ in profile \mathbf{s} . Next, notice that since agent i starts consuming item j at time 0 in \mathbf{s}^* and all other agents use the same strategies in \mathbf{s} and \mathbf{s}^* , it holds that every agent $k \neq i$ starts consuming item j in \mathbf{s}^* no sooner than she does in \mathbf{s} . This means that in profile \mathbf{s}^* , agents other than i will need more time to consume $\beta \cdot \alpha$; in particular they will need at least $\beta t_j(\mathbf{s})$ time, so $t_j(\mathbf{s}^*) \geq \beta t_j(\mathbf{s})$. However, from the previous paragraph we know that they will consume at least $\frac{1}{2} - \frac{1}{4}t_j(\mathbf{s})$, so letting $\beta = \frac{1}{\alpha} (\frac{1}{2} - \frac{1}{4}t_j(\mathbf{s}))$ we get

$$\begin{aligned} t_j(\mathbf{s}^*) &\geq \beta t_j(\mathbf{s}) \geq t_j(\mathbf{s}) \left(\frac{1}{2} - \frac{1}{4} \cdot t_j(\mathbf{s}) \right) \frac{1}{\alpha} \\ &\geq t_j(\mathbf{s}) \left(\frac{1}{2} - \frac{1}{4} \cdot t_j(\mathbf{s}) \right) \geq \frac{1}{4} \cdot t_j(\mathbf{s}) \end{aligned}$$

□

Now we can lower bound the utility of an agent at any pure Nash equilibrium.

Lemma 2. *Let \mathbf{u} be the profile of true agent valuations and let \mathbf{s} be a pure Nash equilibrium. For any agent i and any item j it holds that the utility of agent i at \mathbf{s} is at least $\frac{1}{4} \cdot t_j(\mathbf{s}) \cdot u_{ij}$.*

Proof. Let $\mathbf{s}' = (s'_i, \mathbf{s}_{-i})$ be the strategy profile obtained from \mathbf{s} when agent i deviates to the strategy s'_i where s'_i is some strategy such that $j \succ_{s'_i} \ell$ for all items $\ell \neq j$. Since \mathbf{s} is a pure Nash equilibrium, it holds that $u_i(PS_i(\mathbf{s})) \geq u_i(PS_i(\mathbf{s}')) \geq t_j(\mathbf{s}') \cdot u_{ij}$, where the last inequality holds since the utility of agent i is at least as much as the utility she obtains from the consumption of item j . By Lemma 1, it holds that $t_j(\mathbf{s}') \geq \frac{1}{4} \cdot t_j(\mathbf{s})$ and hence $u_i(PS_i(\mathbf{s})) \geq \frac{1}{4} \cdot t_j(\mathbf{s}) \cdot u_{ij}$. □

We can now prove the pure Price of Anarchy guarantee of the mechanism.

Theorem 1. *The pure Price of Anarchy of Probabilistic Serial is $O(\sqrt{n})$.*

Proof. Let \mathbf{u} be any profile of true agents' valuations and let \mathbf{s} be any pure Nash equilibrium. First, note that by reporting truthfully, every agent i can get an allocation that is at least as good as $(1/n, \dots, 1/n)$, regardless of other agents' strategies. To see this, first consider time $t = 1/n$ and observe that during the interval $[0, 1/n]$, agent i is consuming her favorite item (say a_1) and hence $p_{ia_1} \geq 1/n$. Next, consider time $\tau = 2/n$ and observe that during the interval $[0, 2/n]$, agent i is consuming one or both of her two favorite items (a_1 and a_2) and hence $p_{ia_1} + p_{ia_2} \geq 2/n$. By a similar argument, for any k , it holds that

$\sum_{j=1}^n p_{ia_j} \geq k/n$. This implies that regardless of other agents' strategies, agent i can achieve a utility of at least $\frac{1}{n} \sum_{j=1}^n u_{ij}$. Since \mathbf{s} is a pure Nash equilibrium, it holds that $u_i(PS_i(\mathbf{s})) \geq (1/n) \sum_{j=1}^n u_{ij}$ as well. Summing over all agents, we get that $SW_{PS}(\mathbf{u}, \mathbf{s}) \geq (1/n) \sum_{i=1}^n \sum_{j=1}^n u_{ij} = 1$. If $SW_{OPT}(\mathbf{u}) \leq \sqrt{n}$ then we are done, so assume $SW_{OPT}(\mathbf{u}) > \sqrt{n}$.

Because PS is neutral we can assume $t_j(\mathbf{s}) \leq t_{j'}(\mathbf{s})$ for $j < j'$ without loss of generality. Observe that for all $j = 1, \dots, n$, it holds that $t_j(\mathbf{s}) \geq j/n$. This is true because for any $t \in [0, 1]$, by time t , exactly tn mass of items must have been consumed by the agents. Since j is the j th item that is entirely consumed, by time $t_j(\mathbf{s})$, the mass of items that must have been consumed is at least j . By this, we get that $t_j(\mathbf{s}) \cdot n \geq j$, which implies $t_j(\mathbf{s}) \geq j/n$.

For each j let i_j be the agent that gets item j in the optimal allocation and for ease of notation, let w_{i_j} be her valuation for the item. Now by Lemma 2, it holds that

$$u_{i_j}(PS(\mathbf{s})) \geq \frac{1}{4} \frac{j}{n} w_{i_j} \quad \text{and} \quad SW_{PS}(\mathbf{u}, \mathbf{s}) \geq \frac{1}{4} \sum_{j=1}^n \frac{j}{n} w_{i_j}.$$

The Price of Anarchy is then at most

$$\frac{4 \sum_{j=1}^n w_{i_j}}{\sum_{j=1}^n j \cdot w_{i_j} / n}.$$

Consider the case when the above ratio is maximized and let k be an integer such that $k \leq \sum_{j=1}^n w_{i_j} \leq k + 1$. Then it must be that $w_{i_j} = 1$ for $j = 1, \dots, k$ and $w_{i_j} = 0$, for $k+2 \leq i_j \leq n$. Hence the maximum ratio is $(k + w_{i_{k+1}})/(aw_{i_{k+1}} + b)$, for some $a, b > 0$, which is monotone for $w_{i_{k+1}}$ in $[0, 1]$. Therefore, the maximum value of $(k + w_{i_{k+1}})/(aw_{i_{k+1}} + b)$ is achieved when either $w_{i_{k+1}} = 0$ or $w_{i_{k+1}} = 1$. As a result, the maximum value of the ratio is obtained when $\sum_{i=1}^n w_{i_{k+1}} = k$ for some k . By simple calculations, the Price of Anarchy should be at most:

$$\frac{4k}{\sum_{j=1}^k \frac{j}{n}} \leq \frac{4k}{\frac{k(k-1)}{2n}} = \frac{8n}{k-1},$$

so the Price of Anarchy is maximized when k is minimized. By the argument earlier, $k > \sqrt{n}$ and hence the ratio is $O(\sqrt{n})$. \square

In Section 5, we extend Theorem 1 to broader solution concepts and the incomplete information setting.

Random Priority

We also consider another very well-known mechanism, Random Priority, often referred to as Random Serial Dictatorship. The mechanism first fixes an ordering of the agents uniformly at random and then according to that ordering, it sequentially matches them with their most preferred item that is still available.

Filos-Ratsikas et al. [15] proved that the welfare in any truthtelling equilibrium is an $\Omega(1/\sqrt{n})$ -fraction of the maximum social welfare. While Random Priority has other equilibria as well, to establish the Price of Anarchy bound, it suffices to observe that at least for distinct valuations, any strategy other than truthtelling does not affect the allocation and hence it does not affect the social welfare. Intuitively, since agents pick their most preferred items, any equilibrium strategy would place the most preferred available items on top of the preference list, while the ordering of the items that are not picked does not affect the allocation of other agents. For valuations that are not distinct, the argument can be adapted using small perturbations of the values, losing only a small fraction of welfare.

We first we prove the following lemma.

Lemma 3. *If valuations are distinct, the social welfare is the same in all mixed Nash equilibria of Random Priority.*

Proof. Let i be an agent, and let B be a subset of the items. Let \mathbf{s} be a mixed Nash equilibrium with the property that with positive probability, i will be chosen to select an item at a point when B is the set of remaining items. In that case (by distinctness of i 's values), i 's strategy should place agent i 's favourite item in B on the top of the preference list among items in B . Suppose that for items j and j' , there is no set of items B that may be offered to i with positive probability, in which either j or j' is optimal. Then i may rank them either way, i.e. can announce $j \succ_i j'$ or $j' \succ_i j$. However, that choice has no effect on the other agents, in particular it cannot affect their social welfare. \square

Given the main theorem in [15], Lemma 3 implies the following.

Corollary 1. *If valuations are distinct, the Price of Anarchy of Random Priority is $\Theta(\sqrt{n})$.*

The same guarantee on the Price of Anarchy holds even when the true valuations of agents are not necessarily distinct.

Theorem 2. *The Price of Anarchy of Random Priority is $O(\sqrt{n})$.*

Proof. We know from [15] that the social welfare of Random Priority given truthful reports, is within $O(\sqrt{n})$ of the social optimum. The social welfare of a (mixed) Nash equilibrium \mathbf{q} cannot be worse than the worst pure profile from \mathbf{q} that occurs with positive probability, so let \mathbf{s} be such a pure profile. We will say that agent i *misranks* items j and j' if $j \succ_i j'$, but $j' \succ_{s_i} j$.

If an agent misranks two items for which she has distinct values, it is because she has 0 probability in \mathbf{s} to receive either item. So we can change \mathbf{s} so that no items are misranked, without affecting the social welfare or the allocation. For items that the agent values equally (which are then not misranked) we can apply arbitrarily small perturbations to make them distinct. Profile \mathbf{s} is thus consistent with rankings of items according to perturbed values and is truthful with respect to these values, which, being arbitrarily close to the true ones, have optimum social welfare arbitrarily close to the true optimal social welfare. \square

Theorem 2 can be extended to solution concepts more general than the mixed Nash equilibrium. Again, the details are included in Section 5.

4 Lower Bounds

In this section, we prove our main lower bound. Note that the result holds for any mechanism, including randomized and cardinal mechanisms. Since we are interested in mechanisms with good properties, it is natural to consider those mechanisms that have pure Nash equilibria.

Theorem 3. *The pure Price of Anarchy of any mechanism is $\Omega(\sqrt{n})$.*

Proof. Let $n = k^2$ for some $k \in \mathbb{N}$. Let M be a mechanism and consider the following valuation profile \mathbf{u} . There are \sqrt{n} sets of agents and let G_j denote the j -th set. For every $j \in \{1, \dots, \sqrt{n}\}$ and every agent $i \in G_j$, let $u_{ij} = 1/n + \alpha$ and $u_{ik} = 1/n - \alpha/(n-1)$, for $k \neq j$, where α is sufficiently small. Let \mathbf{s} be a pure Nash equilibrium and for every set G_j , let $i_j = \arg \min_{i \in G_j} p_{ij}^{M, \mathbf{s}}$ (break ties arbitrarily). Observe that for all $j = 1, \dots, \sqrt{n}$, it holds that $p_{i_j j}^{M, \mathbf{s}} \leq 1/\sqrt{n}$ and let $I = \{i_1, i_2, \dots, i_{\sqrt{n}}\}$. Now consider the valuation profile \mathbf{u}' where:

- For every agent $i \notin I$, $u'_i = u_i$.
- For every agent $i_j \in I$, let $u'_{i_j j} = 1$ and $u'_{i_j k} = 0$ for all $k \neq j$.

We claim that \mathbf{s} is a pure Nash equilibrium under \mathbf{u}' as well. For agents not in I , the valuations have not changed and hence they have no incentive to deviate. Assume now for contradiction that some agent $i \in I$ whose most preferred item is item j could deviate to some beneficial strategy s'_i . Since agent i only values item j , this would imply that $p_{ij}^{M, (s'_i, \mathbf{s}-i)} > p_{ij}^{M, \mathbf{s}}$. However, since agent i values all items other than j equally under u_i and her most preferred item is item j , such a deviation would also be beneficial under profile \mathbf{u} , contradicting the fact that \mathbf{s} is a pure Nash equilibrium.

Now consider the expected social welfare of M under valuation profile \mathbf{u}' at the pure Nash equilibrium \mathbf{s} . For agents not in I and taking α to be less than $1/n^3$, the contribution to the social welfare is at most 1. For agents in I , the contribution to the welfare is then at most $(1/\sqrt{n})\sqrt{n} + 1$ and hence the expected social welfare of M is at most 3. As the optimal social welfare is at least \sqrt{n} , the bound follows. \square

Interestingly, if we restrict our attention to *deterministic* mechanisms, then we can prove that only trivial pure Price of Anarchy guarantees are achievable.

Theorem 4. *The pure Price of Anarchy of any deterministic mechanism is $\Omega(n^2)$.*

Proof. Let M be a deterministic mechanism that always has a pure Nash equilibrium. Let \mathbf{u} be a valuation profile such that for all agents i and i' , it holds that $u_i = u_{i'}$, $u_{i1} = 1/n + 1/n^3$ and $u_{ij} > u_{ik}$ for $j < k$. Let \mathbf{s} be a pure Nash equilibrium for this profile and assume without loss of generality that $M_i(\mathbf{s}) = i$.

Now fix another true valuation profile \mathbf{u}' such that $u'_1 = u_1$ and for agents $i = 2, \dots, n$, $u'_{i, i-1} = 1 - \epsilon'_{i, i-1}$ and $u'_{ij} = \epsilon'_{ij}$ for $j \neq i-1$, where $0 \leq \epsilon'_{ij} \leq 1/n^3$,

$\sum_{j \neq i-1} \epsilon'_{ij} = \epsilon'_{i,i-1}$ and $\epsilon'_{ij} > \epsilon'_{ik}$ if $j < k$ when $j, k \neq i-1$. Intuitively, in profile \mathbf{u}' , each agent $i \in \{2, \dots, n\}$ has valuation close to 1 for item $i-1$ and small valuations for all other items. Furthermore, she prefers items with smaller indices, except for item $i-1$.

We claim that \mathbf{s} is a pure Nash equilibrium under true valuation profile \mathbf{u} as well. Assume for contradiction that some agent i has a benefiting deviation, which matches her with an item that she prefers more than i . But then, since the set of items that she prefers more than i in both \mathbf{u} and \mathbf{u}' is $\{1, \dots, i\}$, the same deviation would match her with a more preferred item under \mathbf{u} as well, contradicting the fact that \mathbf{s} is a pure Nash equilibrium. It holds that $SW_{OPT}(\mathbf{u}') \geq n-2$ whereas the social welfare of M is at most $2/n$ and the theorem follows. \square

The mechanism that naively maximizes the sum of the reported valuations with no regard to incentives, when equipped with a lexicographic tie-breaking rule has pure Nash equilibria and also achieves the above ratio in the worst-case, which means that the bounds are tight.

5 General Solution Concepts

In the previous sections, we employed the pure Nash equilibrium as the solution concept for bounding the inefficiency of mechanisms, mainly because of its simplicity. Here, we describe how to extend our results to broader well-known equilibrium concepts in the literature. Due to lack of space, we will only discuss the two most general solution concepts, the *coarse correlated equilibrium* for complete information and the *Bayes-Nash equilibrium* for incomplete information. Since other concepts (like the mixed-Nash equilibrium for instance) are special cases of those two, it suffices to use those for our extensions.

Definition 2. *Given a mechanism M , let \mathbf{q} be a distribution over strategies. Also, for any distribution Δ let Δ_{-i} denote the marginal distribution without the i th index. Then a strategy profile \mathbf{q} is called a*

1. coarse correlated equilibrium if

$$\mathbb{E}_{\mathbf{s} \sim \mathbf{q}} [u_i(M_i(\mathbf{s}))] \geq \mathbb{E}_{\mathbf{s} \sim \mathbf{q}} [u_i(M_i(s'_i, \mathbf{s}_{-i}))],$$

2. Bayes-Nash equilibrium for a distribution $\Delta_{\mathbf{u}}$ where each $(\Delta_{\mathbf{u}})_i$ is independent, if when $\mathbf{u} \sim \Delta_{\mathbf{u}}$ then $\mathbf{q}(\mathbf{u}) = \times_i q_i(u_i)$ and for all u_i in the support of $(\Delta_{\mathbf{u}})_i$,

$$\mathbb{E}_{\mathbf{u}_{-i}, \mathbf{s} \sim \mathbf{q}(\mathbf{u})} [u_i(M_i(\mathbf{s}))] \geq \mathbb{E}_{\mathbf{u}_{-i}, \mathbf{s}_{-i} \sim \mathbf{q}_{-i}(\mathbf{u}_{-i})} [u_i(M_i(s'_i, \mathbf{s}_{-i}))]$$

where the given inequalities hold for all agents i , and (pure) deviating strategies s'_i . Also notice that for randomized mechanisms definitions are with respect to an expectation over the random choices of the mechanism.

The coarse correlated and the Bayesian Price of Anarchy are defined similarly to the pure Price of Anarchy.

Again, first we mention that we can obtain the extensions to Random Priority rather straightforwardly, based on the fact that even when using probability mixtures over strategies, an agent will always (in every realization) pick her most preferred item among the set of available items when she is chosen. In other words, any pure strategy in the support of the distribution will rank the most preferred available item first, and the ordering of the remaining items does not affect the distribution.

Theorem 5. *The coarse correlated Price of Anarchy of Random Priority is $O(\sqrt{n})$. The Bayesian Price of Anarchy of Random Priority is $O(\sqrt{n})$.*

Proof. For the correlated Price of Anarchy, the argument is very similar to the one used in the proof of Theorem 2. Again, if any strategy in the support of a correlated equilibrium \mathbf{q} misranks two items j and j' for any agent i , it can only be because agent i has 0 probability of receiving those items, otherwise agent i would deviate to truth-telling, violating the equilibrium condition. The remaining steps are exactly the same as in the proof of Theorem 2.

For the incomplete information case, consider any Bayes-Nash equilibrium $\mathbf{q}(\mathbf{u})$ and let \mathbf{u} be a any sampled valuation profile. The expected social welfare of the Random Priority can be written as $\mathbb{E}_{\mathbf{u}} [\mathbb{E}_{s \sim \mathbf{q}(\mathbf{u})} [u_i(\mathbf{s})]]$. Using the same argument as the one in the proof of Theorem 2, we can lower bound the quantity $\mathbb{E}_{s \sim \mathbf{q}(\mathbf{u})} [u_i(\mathbf{s})]$ by $\Omega\left(\frac{SW_{OPT}(\mathbf{u})}{\sqrt{n}}\right)$ and the bound follows. \square

Next, we turn to Probabilistic Serial and prove the Price of Anarchy guarantees, with respect to coarse correlated equilibria and Bayes-Nash equilibria. Before we state our theorems however, we will briefly discuss the connection of those extensions with the *smoothness* framework of Roughgarden [25]. According to the definition in [25], a game is (λ, μ) -smooth if it satisfies the following condition

$$\sum_{i=1}^n u_i(s_i^*, \mathbf{s}_{-i}) \geq \lambda SW(\mathbf{s}^*) - \mu SW(\mathbf{s}), \quad (1)$$

where \mathbf{s}^* is a pure strategy profile that corresponds to the optimal allocation and \mathbf{s} is any pure strategy profile. It is not hard to see that a (λ, μ) -smooth game has a Price of Anarchy bounded by $(\mu + 1)/\lambda$.

Since establishing that a game is smooth also implies a pure Price of Anarchy bound, an alternative way of attempting to prove Theorem 1 would be to try to show smoothness of the game induced by PS, for $\mu/\lambda = \sqrt{n}$. However, this seems to be a harder task than what we actually do, since in such a proof, one would have to argue about the utilities of agents and possibly reason about the relative preferences for other items, other than the item they are matched with in the optimal allocation. Our approach only needs to consider those items, and hence it seems to be simpler.

An added benefit to the smoothness framework is the existence of the *extension theorem* in [25], which states that for a (λ, μ) -smooth game, the Price

of Anarchy guarantee extends to broader solution concepts verbatim, without any extra work. At first glance, one might think that proving smoothness for the game induced by PS might be worth the extra effort, since we would get the extensions “for free”. A closer look at our proofs however shows that our approach is very similar to the proof of the extension theorem but using an alternative, simpler condition.

Specifically, the analysis in [25] uses Inequality 1 as a building block and substitutes the inequality into the expectations that naturally appear when considering randomized strategies. This can be done because the condition applies to all strategy profiles \mathbf{s} , when \mathbf{s}^* is an optimal strategy profile. This is exactly what we do as well, but we use the inequality $t_j(s_i^*, \mathbf{s}_{-i}) \geq \frac{1}{4} \cdot t_j(\mathbf{s})$ instead, which is simpler but sufficient since it only applies to the game at hand. If $OPT_i(\mathbf{u}) = j$, which is what we use in the proof of Theorem 1, then (s_i^*, \mathbf{s}_{-i}) can be thought of as a profile where an agent deviates to her strategy in the optimal profile and hence the left-hand side of the inequality is analogous to the left-hand side of Inequality 1. In a sense, the inequality $t_j(s_i^*, \mathbf{s}_{-i}) \geq \frac{1}{4} \cdot t_j(\mathbf{s})$, can be viewed as a “smoothness equivalent” for the game induced by PS, which then allows us to extend the results to broader solution concepts.

First, we extend Theorem 1 to the case where the solution concept is the coarse correlated equilibrium.

Theorem 6. *The coarse correlated Price of Anarchy of Probabilistic Serial is $O(\sqrt{n})$.*

Proof. Let \mathbf{u} be any valuation profile and let i be any agent. Furthermore, let $j = OPT_i(\mathbf{u})$ and let s'_i be the pure strategy that places item j on top of agent i 's preference list. By Lemma 1, the inequality $t_j(s'_i, \mathbf{s}_{-i}) \geq \frac{1}{4}t_j(\mathbf{s})$ holds for every strategy profile \mathbf{s} . In particular, it holds for any pure strategy profile \mathbf{s} where s_i is in the support of the distribution of the mixed strategy q_i of agent i , for any coarse correlated equilibrium \mathbf{q} . This implies that

$$\begin{aligned} \mathbb{E}_{\mathbf{s} \sim \mathbf{q}} [u_i(PS_i(\mathbf{s}))] &\geq \mathbb{E}_{\mathbf{s} \sim \mathbf{q}} [u_i(PS_i(s'_i, \mathbf{s}_{-i}))] \\ &\geq \mathbb{E}_{\mathbf{s} \sim \mathbf{q}} [u_{ij}t_j(s'_i, \mathbf{s}_{-i})] \geq \frac{1}{4}u_{ij}t_j(\mathbf{s}). \end{aligned}$$

where the last inequality holds by Lemma 1. Using this, we can use very similar arguments to the arguments of the proof of Theorem 1 and obtain the bound. \square

For the incomplete information setting, when valuations are drawn from some publically known distributions, we can prove the same upper bound on the Bayesian Price of Anarchy of the mechanism.

Theorem 7. *The Bayesian Price of Anarchy of Probabilistic Serial is $O(\sqrt{n})$.*

Proof. The proof is again similar to the proof of Theorem 1. Let \mathbf{u} be a valuation profile drawn from some distribution satisfying the unit-sum constraint. Let i be any agent and let $j_u = OPT_i(\mathbf{u})$, $i \in [n]$. Note that by a similar argument

as the one used in the proof of Theorem 1, the expected social welfare of PS is at least 1 and hence we can assume that $\mathbb{E}_{\mathbf{u}}[SW_{OPT}(\mathbf{u})] \geq 2\sqrt{2n} + 1$. Observe that in any Bayes-Nash equilibrium $\mathbf{q}(\mathbf{u})$ it holds that

$$\begin{aligned}
\mathbb{E}_{\substack{\mathbf{u} \\ \mathbf{s} \sim \mathbf{q}(\mathbf{u})}} [u_i(\mathbf{s})] &= \mathbb{E}_{u_i} \left[\mathbb{E}_{\substack{\mathbf{u}_{-i} \\ \mathbf{s} \sim \mathbf{q}(\mathbf{u})}} [u_i(\mathbf{s})] \right] \\
&\geq \mathbb{E}_{u_i} \left[\mathbb{E}_{\substack{\mathbf{u}_{-i} \\ \mathbf{s}_{-i} \sim \mathbf{q}_{-i}(\mathbf{u}_{-i})}} [u_i(s'_i, \mathbf{s}_{-i})] \right] \\
&\geq \mathbb{E}_{u_i} \left[\mathbb{E}_{\substack{\mathbf{u}_{-i} \\ \mathbf{s}_{-i} \sim \mathbf{q}_{-i}(\mathbf{u}_{-i})}} [u_{ij_u} t_{j_u}(s'_i, \mathbf{s}_{-i})] \right] \\
&\geq \mathbb{E}_{u_i} \left[\mathbb{E}_{\substack{\mathbf{u}_{-i} \\ \mathbf{s} \sim \mathbf{q}(\mathbf{u})}} \left[\frac{1}{4} u_{ij_u} t_{j_u}(\mathbf{s}) \right] \right] \\
&= \frac{1}{4} \mathbb{E}_{\substack{\mathbf{u} \\ \mathbf{s} \sim \mathbf{q}(\mathbf{u})}} [u_{ij_u} t_{j_u}(\mathbf{s})]
\end{aligned}$$

where the last inequality holds by Lemma 1 since s'_i denotes the strategy that puts item j_u on top of agent i 's preference list. Note that this can be a different strategy for every different \mathbf{u} that we sample. For notational convenience, we use s'_i to denote every such strategy. The expected social welfare at the Bayes-Nash equilibrium is then at least

$$\begin{aligned}
\sum_{i=1}^n \mathbb{E}_{\mathbf{u}, \mathbf{s} \sim \mathbf{q}(\mathbf{u})} [u_i(\mathbf{s})] &\geq \frac{1}{4} \sum_{i \in [n]} \mathbb{E}_{\substack{\mathbf{u} \\ \mathbf{s} \sim \mathbf{q}(\mathbf{u})}} [u_{ij_u} t_{j_u}(\mathbf{s})] \\
&\geq \mathbb{E}_{\substack{\mathbf{u} \\ \mathbf{s} \sim \mathbf{q}(\mathbf{u})}} \left[\sum_{i \in [n]} \frac{i}{4n} u_{ij_u} \right] \\
&\geq \mathbb{E}_{\substack{\mathbf{u} \\ \mathbf{s} \sim \mathbf{q}(\mathbf{u})}} \left[\frac{SW_{OPT}(\mathbf{u})(SW_{OPT}(\mathbf{u}) - 1)}{8n} \right] \\
&= \mathbb{E}_{\mathbf{u}} \left[\frac{SW_{OPT}(\mathbf{u})(SW_{OPT}(\mathbf{u}) - 1)}{8n} \right] \\
&\geq \frac{\mathbb{E}_{\mathbf{u}} [(SW_{OPT}(\mathbf{u}))^2] - \mathbb{E}_{\mathbf{u}} [SW_{OPT}(\mathbf{u})]}{8n} \\
&\geq \frac{\mathbb{E}_{\mathbf{u}} [SW_{OPT}(\mathbf{u})]}{2\sqrt{2n}},
\end{aligned}$$

and the bound follows. \square

6 Extensions

6.1 Price of Stability

Theorem 3 bounds the Price of Anarchy of all mechanisms. A more optimistic (and hence stronger when proving lower bounds) measure of efficiency is the *Price of Stability*, i.e. the worst-case ratio over all valuation profiles of the optimal social welfare over the welfare attained at the *best* equilibrium.

We extend Theorem 3 to the Price of Stability of all mechanisms that satisfy a “proportionality” property.

Let $a_1 \succ_i a_2 \succ_i \dots \succ_i a_n$ be the (possibly weak) preference ordering of agent i . A random assignment vector p_i for agent i *stochastically dominates* another random assignment vector q_i if $\sum_{j=1}^k p_{ia_j} \geq \sum_{j=1}^k q_{ia_j}$, for all $k = 1, 2, \dots, n$. The notation that we will use for this relation is $p_i \succ_i^{sd} q_i$.

Definition 3 (Safe strategy). *Let M be a mechanism. A strategy s_i is a safe strategy if for any strategy profile s_{-i} of the other players, it holds that $M_i(s_i, s_{-i}) \succ_i^{sd} (\frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n})$.*

We will say that a mechanism M has a safe strategy if every agent i has a safe strategy s_i in M . We now state our lower bound.

Theorem 8. *The pure Price of Stability of any mechanism that has a safe strategy is $\Omega(\sqrt{n})$.*

Proof. Let M be a mechanism and let $I = \{k + 1, \dots, n\}$ be a subset of agents. Let \mathbf{u} be the following valuation profile.

- For all agents $i \in I$, let $u_{ij} = \frac{1}{k}$ for $j = 1, \dots, k$ and $u_{ij} = 0$ otherwise.
- For all agents $i \notin I$, let $u_{ii} = 1$ and $u_{ij} = 0, j \neq i$.

Now let \mathbf{s} be a pure Nash equilibrium on profile \mathbf{u} and let s'_i be a safe strategy of agent i . The expected utility of each agent $i \in I$ in the pure Nash equilibrium \mathbf{s} is

$$\begin{aligned} \mathbb{E}[u_i(\mathbf{s})] &= \sum_{j \in [n]} p_{ij}(s_i, \mathbf{s}_{-i}) v_{ij} \geq \sum_{j \in [n]} p_{ij}(s'_i, \mathbf{s}_{-i}) v_{ij} \\ &\geq \frac{1}{n} \sum_{j \in [n]} v_{ij} = \frac{1}{n}, \end{aligned}$$

due to the fact that \mathbf{s} is pure Nash equilibrium and s'_i is a safe strategy of agent i . On the other hand, the utility of agent $i \in I$ can be calculated by $\mathbb{E}[u_i(\mathbf{s})] = \sum_{j \in [n]} p_{ij}(s_i, \mathbf{s}_{-i}) v_{ij} = (\sum_{j=1}^k p_{ij})/k$. Because \mathbf{s} is a pure Nash equilibrium, it holds that $\mathbb{E}[u_i] \geq 1/n$, so we get that $\sum_{j=1}^k p_{ij} \geq k/n$ for all $i \in I$. As for the rest of the agents,

$$\sum_{i \in N \setminus I} \sum_{j=1}^k p_{ij} = k - \sum_{i \in I} \sum_{j=1}^k p_{ij} \leq k - (n - k) \frac{k}{n} = \frac{k^2}{n}.$$

This implies that the contribution to the social welfare from agents not in I is at most k^2/n and the expected social welfare of M will be at most $1 + (k^2/n)$. It holds that $SW_{OPT}(\mathbf{u}) \geq k$ and the bound follows by letting $k = \sqrt{n}$. \square

Due to Theorem 8, in order to obtain an $\Omega(\sqrt{n})$ bound for a mechanism M , it suffices to prove that M has a safe strategy. In fact, most reasonable mechanisms, including Random Priority and Probabilistic Serial, as well as all ordinal *envy-free* mechanisms satisfy this property.

Definition 4 (Envy-freeness). *A mechanism M is (ex-ante) envy-free if for all agents i and r and all profiles \mathbf{s} , it holds that $\sum_{j=1}^n p_{ij} s_{ij} \geq \sum_{j=1}^n p_{rj} s_{rj}$. Furthermore, if M is ordinal, then this implies $p_i^{M,\mathbf{s}} \succ_{s_i}^{sd} p_r^{M,\mathbf{s}}$.*

Given the interpretation of a truth-telling safe strategy as a “proportionality” property, the next lemma is not surprising.

Lemma 4. *Let M be an ordinal, envy-free mechanism. Then for any agent i , the truth-telling strategy u_i is a safe strategy.*

Proof. Let $\mathbf{s} = (u_i, \mathbf{s}_{-i})$ be the strategy profile in which agent i is truth-telling and the rest of the agents are playing some strategies \mathbf{s}_{-i} . Since M is envy-free and ordinal, it holds that $\sum_{j=1}^{\ell} p_{ij}^{\mathbf{s}} \geq \sum_{j=1}^{\ell} p_{rj}^{\mathbf{s}}$ for all agents $r \in \{1, \dots, n\}$ and all $\ell \in \{1, \dots, n\}$. Summing up these inequalities for agents $r = 1, 2, \dots, n$ we obtain

$$n \sum_{j=1}^{\ell} p_{ij}^{\mathbf{s}} \geq \sum_{j=1}^{\ell} \sum_{r=1}^n p_{rj}^{\mathbf{s}} = \ell,$$

which implies that $\sum_{j=1}^{\ell} p_{ij}^{\mathbf{s}} \geq \frac{\ell}{n}$, for all $i \in \{1, \dots, n\}$, and for all $\ell \in \{1, \dots, n\}$. \square

Note that since Probabilistic Serial is ordinal and envy-free [9], by Lemma 4, it has a safe strategy and hence Theorem 8 applies. It is not hard to see that Random Priority has a safe strategy too.

Lemma 5. *Random Priority has a safe strategy.*

Proof. Since Random Priority first fixes an ordering of agents uniformly at random, every agent i has a probability of $1/n$ to be selected first to choose an item, a probability of $2/n$ to be selected first or second and so on. If the agent ranks her items truthfully, then for every $\ell = 1, \dots, n$, it holds that $\sum_{i=1}^{\ell} p_{ij} \geq \ell/n$. \square

In a sense, the safe strategy property is essential for the bound to hold; one can show that the *randomly dictatorial* mechanism, that matches a uniformly chosen agent with her most preferred item and the rest of the agents with items based solely on that agent’s reports achieves a constant Price of Stability. On the other hand, the Price of Anarchy of the mechanism is $\Omega(n)$. It would be interesting to show whether Price of Anarchy guarantees imply Price of Stability lower bounds in general.

6.2 Unit-range representation

Our second extension is concerned with the other normalization that is also common in the literature [2,15,28], the unit-range representation, that is, $\max_j u_i(j) = 1$ and $\min_j u_i(j) = 0$. First, the Price of Anarchy guarantees from Section 3 extend directly to the unit-range case. For Random Priority, since the results in [15] hold for this normalization as well, we can apply the same techniques to prove the bounds. For Probabilistic Serial, first, observe that Lemma 2 holds independently of the representation. Secondly, in the proof of Theorem 1, it now holds that

$$SW_{PS}(\mathbf{u}, \mathbf{s}) \geq \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n u_{ij} \geq 1,$$

which is sufficient for bounding the Price of Anarchy when $SW_{OPT}(\mathbf{u}) \leq \sqrt{n}$. Finally, the arguments for the case when $SW_{OPT}(\mathbf{u}) \leq \sqrt{n}$ hold for both representations.

Next, we present a Price of Anarchy lower bound for deterministic mechanisms. First, we prove the following lemma about the structure of equilibria of deterministic mechanisms. Note that the lemma holds independently of the choice of representation.

Lemma 6. *The set of pure Nash equilibria of any deterministic mechanism is the same for all valuation profiles that induce the same preference orderings of valuations.*

Proof. Let \mathbf{u} and \mathbf{u}' be two different valuation profiles that induce the same preference ordering. Let \mathbf{s} be a pure Nash equilibrium under true valuation profile \mathbf{u} and assume for contradiction that it is not a pure Nash equilibrium under \mathbf{u}' . Then, there exists an agent i who by deviating from \mathbf{s} is matched to a more preferred item according to u'_i . But that item would also be more preferred according to u_i and hence she would have an incentive to deviate from \mathbf{s} under true valuation profile \mathbf{u} , contradicting the fact that \mathbf{s} is a pure Nash equilibrium. \square

Using Lemma 6, we can then prove the following theorem.

Theorem 9. *The Price of Anarchy of any deterministic mechanism that always has pure Nash equilibria is $\Omega(n)$ for the unit-range representation.*

Proof. Let M be a deterministic mechanism that always has a pure Nash equilibrium and let \mathbf{u} be a valuation profile such that for all agents i and i' , it holds that $u_i = u_{i'}$ and $u_{ij} > u_{ik}$, for all items $i < k$. Let \mathbf{s} be a pure Nash equilibrium for this profile and assume without loss of generality that $M_i(\mathbf{s}) = i$. By Lemma 6, \mathbf{s} is a pure Nash equilibrium for any profile \mathbf{u} that induces the above ordering of valuations. In particular, it is a pure Nash equilibrium for a valuation profile satisfying

- For agents $i = 1, \dots, \frac{n}{2}$, $u_{i1} = 1$ and $u_{ij} < \frac{1}{n^3}$, for $j > 1$.

- For agents $i = \frac{n}{2} + 1, \dots, n$, $u_{ij} > 1 - \frac{1}{n^3}$ for $j = 1, \dots, n/2$ and $u_{ij} < \frac{1}{n^3}$ for $j = \frac{n}{2} + 1, \dots, n$.

It holds that $OPT(\mathbf{u}) \geq \frac{n}{2}$, whereas the social welfare of M is at most 2 and the theorem follows. \square

Again, similarly to the corresponding bound in Section 4, the mechanism that naively maximizes the sum of the reported valuations has pure Nash equilibria and achieves the above bound.

More importantly, it is not clear whether the general lower bound on the Price of Anarchy of all mechanisms that we proved in Theorem 3 extends to the unit-range representation as well. In fact, we do not know of any bound for the unit-range case and proving one seems to be a quite complicated task. As a first step in that direction, the following theorem obtains a lower bound for ϵ -approximate (pure) Nash equilibria. A strategy profile is an ϵ -approximate pure Nash equilibrium if no agent can deviate to another strategy and improve her utility by more than ϵ . While the following result applies for any positive ϵ , it is weaker than a corresponding result for exact equilibria.

Theorem 10. *Let M be a mechanism and let $\epsilon \in (0, 1)$. The ϵ -approximate Price of Anarchy of M is $\Omega(n^{1/4})$ for the unit-range representation.*

Proof. Assume $n = k^2$, where $k \in \mathbb{N}$ will be the size of a subset I of “important” agents. We consider valuation profiles where, for some parameter $\delta \in (0, 1)$,

- all agents have value 1 for item 1,
- there is a subset I of agents with $|I| = k$ for which any agent $i \in I$ has value δ^2 for any item $j \in \{2, \dots, k + 1\}$ and 0 for all other items,
- for agent $i \notin I$, i has value δ^3 for items $j \in \{2, \dots, k + 1\}$ and 0 for all other items.

Let \mathbf{u} be such a valuation profile and let \mathbf{s} be a Nash equilibrium. In the optimal allocation members of I receive items $\{2, \dots, k + 1\}$ and such an allocation has social welfare $k\delta^2 + 1$.

First, we claim that there are $k(1 - 2\delta)$ members of I whose payoffs in \mathbf{s} are at most δ ; call this set X . If that were false, then there would be more than $2k\delta$ members of I whose payoffs in \mathbf{s} were more than δ . That would imply that the social welfare of \mathbf{s} was more than $2k\delta^2$, which would contradict the optimal social welfare attainable, for large enough n (specifically, $n > 1/\delta^4$).

Next, we claim that there are at least $k(1 - 2\delta)$ non-members of I whose probability (in \mathbf{s}) to receive any item in $\{1, \dots, k + 1\}$ is at most $4(k + 1)/n$; call this set Y . To see this, observe that there are at least $\frac{3}{4}n$ agents who all have probability $\leq 4/n$ to receive item 1. Furthermore, there are at least $3n/4$ agents who all have probability $\leq 4k/n$ to receive an item from the set $2, \dots, k + 1$. Hence there are at least $n/2$ agents whose probabilities to obtain these items satisfy both properties.

We now consider the operation of swapping the valuations of the agents in sets X and Y so that the members of I from X become non-members, and vice

versa. We will argue that given that they were best-responding beforehand, they are δ -best-responding afterwards. Consequently \mathbf{s} is an δ -NE of the modified set of agents. The optimum social welfare is unchanged by this operation since it only involves exchanging the payoff functions of pairs of agents. We show that the social welfare of \mathbf{s} is some fraction of the optimal social welfare, that goes to 0 as n increases and δ decreases.

Let I' be the set of agents who, after the swap, have the higher utility of δ^2 for getting items from $\{2, \dots, k+1\}$. That is, I' is the set of agents in Y , together with I minus the agents in X .

Following the above valuation swap, the agents in X are δ -best responding. To see this, note that these agents have had a reduction to their utilities for the outcome of receiving items from $\{2, \dots, k+1\}$. This means that a profitable deviation for such agents should result in them being more likely to obtain item 1, in return for them being less likely to obtain an item from $\{2, \dots, k+1\}$. However they cannot have probability more than δ to receive item 1, since that would contradict the property that their expected payoff was at most δ .

After the swap, the agents in Y are also δ -best responding. Again, these agents have had their utilities increased from δ^3 to δ^2 for the outcome of receiving an item from $\{2, \dots, k+1\}$. Hence any profitable deviation for such an agent would involve a reduction in the probability to get item 1 in return for an increased probability to get an item from $\{2, \dots, k+1\}$. However, since the payoff for any item from $\{2, \dots, k+1\}$ is only δ^2 , such a deviation pays less than δ .

Finally, observe that the social welfare of \mathbf{s} under the new profile (after the swap) is at most $1 + 3k\delta^3$. To see this, note that (by an earlier argument and the definition of I') $k(1 - 2\delta)$ members of I' have probability at most $4(k+1)/n$ to receive any item from $\{1, \dots, k+1\}$. To upper bound the expected social welfare, note that item 1 contributes 1 to the social welfare. Items in $\{2, \dots, k+1\}$ contribute in total, δ^2 times the expected number of members of I' who get them, plus δ^3 times the expected number of non-members of I' who get them, which is at most $\delta^2 k 2\delta + \delta^3 k(1 - 2\delta)$ which is less than $3k\delta^3$.

Overall, the price of anarchy is at least $(k\delta^2 + 1)/3k\delta^3$, which is more than $1/\delta$. The statement of the theorem is obtained by choosing δ to be less than ϵ , n large enough for the arguments to hold for the chosen δ , i.e. $n > 1/\delta^4$. \square

7 Conclusion and Future Work

Our results are rather negative: we identify a non-constant lower bound on the Price of Anarchy for one-sided matching, and find a matching upper bound achieved by well-known ordinal mechanisms. However, such negative results are important to understand the challenge faced by a social-welfare maximizer: for example, we establish that it is not enough to elicit cardinal valuations, in order to obtain good social welfare. It may be that better welfare guarantees should use some assumption of truth-bias, or some assumption of additional structure in agents' preferences.

An interesting direction of research would be to identify conditions on the valuation space that allow for constant values of the Price of Anarchy or impose some distributional assumption on the inputs and quantify the average loss in welfare due to selfish behavior. For the general, worst-case setting, one question raised is whether one can obtain Price of Anarchy or Price of Stability bounds that match our upper bounds for the unit-range representation as well.

Acknowledgements

The authors would like to thank Piotr Krysta for useful discussion.

References

1. Atila Abdulkadiroğlu and Tayfun Sönmez. Matching Markets: Theory and Practice. *Advances in Economics and Econometrics (Tenth World Congress)*, pages 3–47, 2013.
2. Marek Adamczyk, Piotr Sankowski, and Qiang Zhang. Efficiency of Truthful and Symmetric Mechanisms in One-sided Matching. In *Proceedings of the 7th Symposium of Algorithmic Game Theory*, pages 13–24. Springer, 2014.
3. Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
4. Haris Aziz, Felix Brandt, and Markus Brill. The computational complexity of random serial dictatorship. *Economics Letters*, 121(3):341–345, 2013.
5. Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Nina Narodytska, and Toby Walsh. Equilibria under the probabilistic serial rule. *arXiv preprint arXiv:1502.04888*, 2015.
6. Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Nina Narodytska, and Toby Walsh. Manipulating the probabilistic serial rule. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1451–1459. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
7. Haris Aziz, Serge Gaspers, Nick Mattei, Nina Narodytska, and Toby Walsh. Strategic aspects of the probabilistic serial rule for the allocation of goods. *arXiv preprint arXiv:1401.6523*, 2014.
8. Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. Social Welfare in One-Sided Matching Markets without Money. In *APPROX-RANDOM*, pages 87–98, 2011.
9. Anna Bogomolnaia and Hervé Moulin. A New Solution to the Random Assignment Problem. *Journal of Economic Theory*, 100:295–328, 2001.
10. Steven J Brams, Michal Feldman, John K Lai, Jamie Morgenstern, and Ariel D Procaccia. On maxsum fair cake divisions. In *AAAI*, 2012.
11. Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. The efficiency of fair division. *Theory of Computing Systems*, 50(4):589–610, 2012.
12. Deeparnab Chakrabarty and Chaitanya Swamy. Welfare maximization and truthfulness in mechanism design with ordinal preferences. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 105–120, 2014.

13. George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian Combinatorial Auctions. In *ICALP '08: Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*. Springer-Verlag, July 2008.
14. Ozgün Ekici and Onur Kesten. On the Ordinal Nash Equilibria of the Probabilistic Serial Mechanism. Technical report, working paper, Tepper School of Business, Carnegie Mellon University, 2010.
15. Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 1–12, 2014.
16. Mingyu Guo and Vincent Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*, pages 881–888, 2010.
17. Tadashi Hashimoto, Daisuke Hirata, Onur Kesten, Morimitsu Kurino, and M Utku Ünver. Two axiomatic approaches to the probabilistic serial mechanism. *Theoretical Economics*, 9(1):253–277, 2014.
18. Aanund Hylland and Richard Zeckhauser. The Efficient Allocation of Individuals to Positions. *The Journal of Political Economy*, 87(2):293–314, 1979.
19. Akshay-Kumar Katta and Jay Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic theory*, 131(1):231–250, 2006.
20. O Kesten. Probabilistic serial and top trading cycles from equal division for the random assignment problem. Technical report, mimeo, 2006.
21. Fuhito Kojima and Mihai Manea. Incentives in the probabilistic serial mechanism. *Journal of Economic Theory*, 145(1):106–123, 2010.
22. Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *STACS 99*, pages 404–413. Springer, 1999.
23. Piotr Krysta, David Manlove, Baharak Rastegari, and Jinshan Zhang. Size versus truthfulness in the House Allocation problem. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 453–470. ACM, 2014.
24. Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 177–186. ACM, 2009.
25. Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 513–522. ACM, 2009.
26. Vasilis Syrgkanis and Eva Tardos. Composable and Efficient Mechanisms. In *STOC '13: Proceedings of the 45th symposium on Theory of Computing*, November 2013.
27. John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior (60th Anniversary Commemorative Edition)*. Princeton university press, 2007.
28. Lin Zhou. On a Conjecture by Gale about One-Sided Matching Problems. *Journal of Economic Theory*, 52:123–135, 1990.