The Social Welfare of One-Sided Matching Mechanisms

GEORGE CHRISTODOULOU, University of Liverpool
ARIS FILOS-RATSIKAS, University of Oxford
SØREN KRISTOFFER STIIL FREDERIKSEN, Aarhus University
PAUL W. GOLDBERG, University of Oxford
JIE ZHANG, University of Oxford
JINSHAN ZHANG, University of Liverpool

We study the fundamental problem of *one-sided matching* where n agents have unrestricted cardinal valuations over n items and the goal is to maximize social welfare, i.e. the total utility of the agents in an assignment. We measure the quality of a matching mechanism by its *Price of Anarchy* and prove lower and upper bounds on the performance of mechanisms with respect to the two dominant representations of valuation functions in the literature, unit-range and unit-sum.

First, we prove that the two most prominent one-sided matching mechanisms, Probabilistic Serial and Random Priority achieve a Price of Anarchy of $O(\sqrt{n})$, with respect to both representations and very general equilibrium notions, as well as the case of incomplete information. We complement this result with a lower bound of $\Omega(\sqrt{n})$ on the Price of Anarchy of any mechanism for the unit-sum representation, which proves that those two mechanisms are asymptotically optimal among *all* mechanisms for the problem, including randomized and cardinal ones. For the unit-range representation, we prove a lower bound of $\Omega(\sqrt{n})$ on the performance of all *truthful* mechanisms, which implies that Random Priority is optimal among all mechanisms in this class, as well as a lower bound of $\Omega(n^{-1/4})$ on the Price of Anarchy of any mechanism with respect to its ϵ-equilibria.

Additionally, we prove that the *Price of Stability* of any *proportional* mechanism is bounded by $\Omega(\sqrt{n})$; most natural matching mechanisms including Random Priority and Probabilistic Serial are proportional. For deterministic mechanisms, we show strong lower bounds for both representations, which imply that randomization is needed for better welfare guarantees to be achievable.

CCS Concepts: •**Theory of computation** → **Algorithmic game theory and mechanism design; Algorithmic game theory;** *Quality of equilibria;*

General Terms: Theory, Algorithms, Economics

Additional Key Words and Phrases: Algorithmic mechanism design, one-sided matching, social welfare, price of anarchy, truthfulness, approximation ratio

ACM Reference Format:
George Christodoulou, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Paul W. Goldberg, Jie Zhang, and Jinshan Zhang, 2016. The Social Welfare of One-Sided Matching Mechanisms. *ACM Trans. Econ. Comp.* 0, 0, Article 0 (2016), 33 pages.
DOI: 0000001.0000001

G. Christodoulou, P.W. Goldberg and Jinshan Zhang were supported by the EPSRC grant EP/K01000X/1. P.W. Goldberg was supported by COST Action IC1205. A. Filos-Ratsikas and Jie Zhang were supported by the ERC Advanced Grant 321171 (ALGAME). A. Filos-Ratsikas and S.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. This paper includes most of the results from [Filos-Ratsikas et al. 2014] and [Christodoulou et al. 2016].
Author's addresses: G. Christodoulou and Jinshan Zhang, Computer Science Department, University of Liverpool; A. Filos-Ratsikas and Paul W. Goldberg and Jie Zhang, Computer Science Department, University of Oxford; S.K.S. Frederiksen, Computer Science Department, Aarhus University.

DOI: 0000001.0000001

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 lies in the intersection of computer science, social choice and mechanism design. The study of the problem has its roots in the seminal papers of Shapley and Scarf [1974] and Hylland and Zeckhauser [1979] and has found numerous applications throughout the years, such as task allocation, student or medical resident placement [Sönmez and Ünver 2011], or more recently, nationwide organ exchange programs [UNOS 1984].

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 [Hylland and Zeckhauser 1979; Bogomolnaia and Moulin 2001]. In game-theoretic terms, these are the agents' von Neumann-Morgenstern utility functions [Von Neumann and Morgenstern 1944] and the associated preferences are often referred to as *cardinal preferences*.

A *mechanism* is a function that elicits the agents' private valuations for the items and maps them 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). The assumption that the agents report their valuations strategically to maximize their utilities is central to the fields of game theory and mechanism design, which suggest tools for preventing or managing such selfish behaviour. In particular, one approach is to try to eliminate incentives for misreporting altogether, by constructing *truthful* mechanisms, i.e. mechanisms where regardless of the choices of the other agents, an agent is better off by revealing his valuation truthfully. A more general approach would be to study the *Nash equilibria* of the induced game of a mechanism, i.e. strategy profiles from which no agent wishes to unilaterally deviate, and design mechanisms with good equilibria.

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* [Koutsoupias and Papadimitriou 1999], 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 distribution over matchings, instead of a single matching. We are interested in the class of *cardinal* mechanisms, which use the actual numerical values of the preferences, and generalize the ordinal mechanisms, which only make use of the orderings induced by the agents' valuations.

Note that our setting involves no monetary transfers and generally falls under the umbrella of *approximate mechanism design without money* [Procaccia and Tennenholtz 2009]. In general settings without money, one has to fix a canonical representation of the valuations. The two predominant assumptions in the literature are the *unit-sum* representation, i.e. each agent has a total value of 1 for all the items or the unit-range representation, i.e. each agent's valuations lie in $[0, 1]$ with the maximum valuation being 1 and the minimum being 0.

Table I. Table of results

Bounds/Mechanisms	PS	RP	Any	Truthful	Proportional
Unit-Sum					
Price of Anarchy	$\Theta(\sqrt{n})$	$\Theta(\sqrt{n})$	$\Omega(\sqrt{n})$	$\Omega(\sqrt{n})$	$\Omega(\sqrt{n})$
Approximation Ratio	$\times$	$\Theta(\sqrt{n})$	$\times$	$\Omega(\sqrt{n})$	$\times$
Price of Stability	$\Theta(\sqrt{n})$	$\Theta(\sqrt{n})$	$\Theta(1)$	$\Theta(1)$	$\Omega(\sqrt{n})$
Unit-Range					
Price of Anarchy	$\Theta(\sqrt{n})$	$\Theta(\sqrt{n})$	$\Omega(n^{1/4})^\dagger$	$\Omega(\sqrt{n})$	$\Omega(n^{1/4})^\dagger$
Approximation Ratio	$\times$	$\Theta(\sqrt{n})$	$\times$	$\Omega(\sqrt{n})$	$\times$
Price of Stability	$O(\sqrt{n})$	$O(\sqrt{n})$	$\Theta(1)$	$\Theta(1)$	$-$

Table summarizing the main results. "O" bounds are approximation guarantees and "Ω" bounds are inapproximability bounds. A $\Theta(d)$ bound for a class of mechanisms means that no mechanism in the class can perform better than $\Omega(d)$ and there is a mechanism in the class whose performance is $O(d)$. "$\times$" indicates that an approximation ratio result is only meaningful for truthful mechanisms, where truth-telling is guaranteed to be an equilibrium. "†" indicates that the result is with respect to ϵ-equilibria, for any $\epsilon > 0$.

1.1. Our results

In Section 3 we bound the inefficiency of the two best-known mechanisms in the matching literature, *Probabilistic Serial* [Bogomolnaia and Moulin 2001] and *Random Priority* [Abdulkadiroğlu and Sönmez 1998]. In particular, for n agents and n items, their Price of Anarchy is $O(\sqrt{n})$ for both representations. The bounds hold for very general solution concepts, such as *coarse-correlated equilibria* as well as *Bayes-Nash* equilibria, when the valuations of the agents are drawn from some probability distribution. In Section 4 we complement these with a *matching* lower bound (i.e. $\Omega(\sqrt{n})$) that applies to *all* cardinal (randomized) mechanisms for the unit-sum representation. As a result, we conclude that these two *ordinal* mechanisms are optimal. These results suggest that in the presence of selfish behaviour, it does not help a welfare maximizer to ask agents to report preferences more expressively than simply rankings over the items.

For unit-range, we prove a similar lower bound of $\Omega(\sqrt{n})$ on the Price of Anarchy of all truthful mechanisms, again establishing that Random Priority is the best mechanism in this class. In fact, we prove a stronger statement, namely that the *approximation ratio* (i.e. a notion similar to the Price of Anarchy that is with respect to only the truth-telling equilibria) of any truthful mechanism is bounded by $\Omega(\sqrt{n})$. For all mechanisms, under no restrictions, we prove a lower bound of $\Omega(n^{1/4})$ on the Price of Anarchy, with respect to their ϵ-equilibria, for $\epsilon > 0$.

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

Finally, in Section 5, we extend our main lower bound to the *Price of Stability*, a more optimistic measure of efficiency [Anshelevich et al. 2008], which strengthens the negative results even further. Specifically, we prove that the performance of any *proportional* mechanism is bounded by $\Omega(\sqrt{n})$.

An overview of our main results can be found in Table I.

1.2. Discussion and related work

The one-sided matching problem has been a focus of attention of much related literature in economics, computer science and artificial intelligence for many years. Here, we will cover some ground on the major classical approaches in economics as well as the more recent literature in computer science, focusing on aspects that are mostly relevant to our goals. For a more comprehensive exposition of the plethora of results on the topic, the interested reader is referred to the surveys of Sönmez and Ünver [2011] and Abdulkadiroglu and Sönmez [2013] and the references therein.

As we mentioned earlier, the one-sided matching problem was first introduced by Shapley and Scarf [1974] in the context of initial endowments and later on Hylland and Zeckhauser [1979] formulated a version of the problem where items are not initially endowed by agents (and termed this "social endowment"). The matching problem with preferences on both sides of the market was considered earlier in the very influential paper of Gale and Shapley [1962]. Over the years, several different one-sided matching mechanisms have been proposed with various desirable properties related to truthfulness, fairness and economic efficiency.

Perhaps the most intuitive solution would be to fix an ordering of agents and let them pick their most-preferred items according to this ordering; the resulting class of mechanisms is known as *serial dictatorships* [Svensson 1999]. To avoid unfairness issues, one could fix the ordering uniformly at random. The resulting mechanism, which presumably dates back to ancient times, is termed random serial dictatorship or *random priority*. Random Priority has been extensively studied in the literature of one-sided matching [Sönmez and Ünver 2005; Abdulkadiroğlu and Sönmez 1998; Krysta et al. 2014; Hosseini et al. 2015; Bhalgat et al. 2011], mainly due to its simplicity, alongside its desirable properties; the mechanism is fair (in the sense of anonymity), truthful, and for every fixed ordering, the resulting allocation is Pareto efficient, i.e. there is no other allocation that could make some agent more satisfied without making some agent less satisfied. On the other hand, it fails to achieve stronger fairness guarantees, like envy-freeness [Foley 1967], a property in which no agent would rather swap her randomized allocation with any other agent.

In an attempt to mend the fairness issues of Random Priority and achieve better efficiency guarantees, Crès and Moulin [2001] introduced the *Probabilistic Serial* mechanism, which was popularized by Bogomolnaia and Moulin [2001] in a paper very central to the field. In Probabilistic Serial, agents consume items continuously with a fixed speed according to their preferences, starting with their most-preferred item and moving on the the next non-depleted item on their preference list. The resulting fractional allocation can be interpreted as a randomized allocation of indivisible items. The merits of the mechanism are its *ordinal efficiency*, a version of efficiency stronger than ex-post Pareto efficiency and the fact that it is envy-free. Note that the mechanism is not truthful[1] and therefore its performance is naturally best evaluated in equilibrium [Ekici and Kesten 2010].

We note here that 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 due to its popularity and extended study within the literature, with related work on characterizations [Hashimoto et al. 2014; Kesten 2006], extensions [Katta and Sethuraman 2006], strategic aspects [Kojima and Manea 2010; Hosseini et al. 2016] and hardness of manipulation [Aziz et al. 2015b]. Somewhat surprisingly, the Nash equilibria of the mechanism were only recently studied. Aziz et al. [2015a] prove that the mechanism has pure Nash equilibria while Ekici and Kesten [2010] study the *ordinal* equilibria of the mechanism and prove that the desirable properties of the mechanism are not necessarily satisfied for those profiles.

[1]Probabilistic Serial satisfies a much weaker condition that was coined *weak truthfulness* in [Bogomolnaia and Moulin 2001].

Notice that both Random Priority and Probabilistic Serial are ordinal; in fact, the matching literature in economics has been dominated by ordinal mechanisms. Interestingly enough though, Hylland and Zeckhauser in their 1979 paper propose a cardinal mechanism, *the pseudo-market mechanism*. The mechanism first endows agents with artificial budgets of unit capacity and then produces a randomized matching in a market-like fashion: items are treated as divisible commodities, prices are announced and agents purchase their most preferred shares at those prices. The process is repeated until supply meets demand, i.e. all items are entirely allocated and all artificial budgets are exhausted.[2] The pseudo-market mechanism is strong in terms of economic efficiency and fairness, but it is not truthful; truthfulness as a desirable property had already been discussed in [Hylland and Zeckhauser 1979]. In this paper, we will be interested in the performance of *all* mechanisms, including cardinal ones, which for instance precludes the option of employing recent characterization results by Mennle and Seuken [2014] for our lower bounds on truthful mechanisms, since they only apply to ordinal mechanisms.

As one can see from the discussion above, the literature in classical economics has been primarily interested in achieving tradeoffs between economic efficiency, truthfulness and fairness and the limitations of such attempts have been considered in [Zhou 1990] and [Bogomolnaia and Moulin 2001] among others. In the computer science literature, the research direction has shifted towards aggregate measures of social efficiency, similarly to what we do in this paper. A recent branch of study considers ordinal measures of efficiency instead of social welfare maximization though, under the assumption that agents' preferences are only expressed through preference orderings over items. Bhalgat et al. [2011] study the approximation ratio of matching mechanisms, when the objective is the maximization of the *ordinal social welfare*, a notion of efficiency that they define based solely on ordinal information. Other measures of efficiency for one-sided matching were also studied in Krysta et al. [2014], where the authors design truthful mechanisms to approximate the size of a maximum cardinally (or maximum agent weight) Pareto-optimal matching and in Chakrabarty and Swamy [2014] 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 [Bogomolnaia and Moulin 2001; Hylland and Zeckhauser 1979; Von Neumann and Morgenstern 1944; Zhou 1990]. Our results actually suggest that (at least for the unit-sum representation), 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.

At the same time as the appearance of the first conference paper [Filos-Ratsikas et al. 2014] that includes some of the results of the present paper, independently Adamzyck et al. [2014] studied truthful mechanisms for social welfare maximization in one-sided matchings when agents have von Neumann-Morgenstern utilities, normalized in the unit interval, but not necessarily unit-sum or unit-range. Their main result on the approximation ratio of random priority can be combined with some additional arguments to obtain one of our intermediate lemmas but our upper bounds, even those that only apply to truthful mechanisms, are more general and in particular imply theirs.

[2]The pseudo-market mechanism is also sometimes referred to as the *CEEI mechanism* [Bogomolnaia and Moulin 2001], where CEEI stands for "competitive equilibrium from equal incomes", the supply-meets-demand outcome of a market where buyers have equal budgets.

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 for settings without money, initiated by [Christodoulou et al. 2008] and studied under a lot of different variants [Bhawalkar and Roughgarden 2011; Christodoulou et al. 2016; Hassidim et al. 2011; Feldman et al. 2013; de Keijzer et al. 2013; Roughgarden 2014]. 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 [Roughgarden 2009] for games and [Syrgkanis and Tardos 2013] for mechanisms. While our results are not proven using the smoothness condition, our extension technique is similar in spirit.

1.3. Canonical Representation

In standard utility theory, von Neumann-Morgenstern utility functions are well-defined up to positive affine transformations. The choice of transformation does not make a difference when arguing about utilities on an individual basis but it is quite important when considering interpersonal objectives, like social welfare maximization. Applying different transformations to the utility functions would result in different agents having inputs of varying "importance". The standard approach in literature is to fix some canonical representation, or *normalization* of utilities and there are two popular approaches, *unit-sum* and *unit-range*.

In the *unit-sum* normalization, an agent has a total value of 1, which implies that each agent would be equally happy when receiving all the items. Intuitively, this normalization means that every 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 normalization has been applied for social welfare maximization in many settings without money including *fair division* and *cake-cutting* [Brams et al. 2012; Caragiannis et al. 2012; Cohler et al. 2011; Bertsimas et al. 2011; Karp et al. 2014; Cole et al. 2013a], *indivisible and divisible item allocation* [Guo and Conitzer 2010; Brânzei et al. 2014; Feldman et al. 2009; Cole et al. 2013b; Han et al. 2011; Cheung 2016] and *voting* [Boutilier et al. 2012; Caragiannis and Procaccia 2011; Caragiannis et al. 2016] among others.

If we use the same affine transformation for all utility functions, agents" utilities are fixed to lie in an interval with the least preferred and the most preferred candidates mapped to the endpoints of the interval respectively. In the *unit-range* representation, the chosen interval is $[0, 1]$; this is equivalent to the "zero-one rule" proposed by Hausman [Hausman 1995] for normalizing and comparing von Neumann-Morgenstern utilities. The unit-sum representation has been used often in the literature as well [Zhou 1990; Barbera 2010; Feige and Tennenholtz 2010; Filos-Ratsikas and Miltersen 2014].

For a more detailed discussion of the two representations, we refer the read to [Filos-Ratsikas 2015]. Finally, let us remark that without any normalization, non-trivial social welfare guarantees are not possible by any mechanism (e.g. see [Anshelevich and Das 2010] for bounds on truthful mechanisms among others).

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 $u_i : A \to \mathbb{R}$ mapping items to real numbers. Valuation functions are considered to be well-defined modulo positive affine transformations, that is, for item $j : j \to \alpha u_i(j) + \beta$

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 either

— ***unit-sum***, that is $\sum_j u_i(j) = 1$, with $u_i(j) \geq 0$ for all i, j or

— ***unit-range***, i.e. $\max_{j \in A} u_i(j) = 1$ and $\min_{j \in A} u_i(j) = 0$.[3]

Equivalently, we can consider valuation functions as *valuation vectors* $u_i = (u_{i1}, u_{i2}, \ldots, u_{in})$ and let V be the set of all valuation vectors of an agent. Let $\mathbf{u} = (u_1, u_2, \ldots, u_n)$ 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 $\mathrm{s} = (s_1, s_2, \ldots, s_n)$ to be a pure strategy profile, where s_i is the *reported* valuation vector of agent i. We will use $\mathrm{s}_{-\mathbf{i}}$ to denote the strategy profile without the ith coordinate and hence $\mathrm{s} = (s_i, \mathrm{s}_{-\mathbf{i}})$ is an alternative way to denote a strategy profile. A *direct revelation mechanism* without money is a function $M : V^n \to O$ mapping *reported* valuation profiles to matchings. For a randomized mechanism, we define M to be a random map $M : V^n \to O$. Let $M_i(\mathrm{s})$ 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 $p_{ij}^{M,\mathbf{s}} = \Pr[M_i(\mathrm{s}) = j]$ and $p_i^{M,\mathbf{s}} = (p_{i1}^{M,\mathbf{s}}, \ldots, p_{in}^{M,\mathbf{s}})$. 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 s is simply $u_i(M_i(\mathrm{s}))$. For randomized mechanisms, an agent's utility is $\mathbb{E}[u_i(M_i(\mathrm{s}))] = \sum_{j=1}^{n} p_{ij}^{M,\mathbf{s}} u_{ij}$.

Note that agents can be or appear to be indifferent between items and hence both the valuations and the strategies could exhibit *ties*, i.e. for two items j and j', both $v_i(j) = v_i(j')$ and $s_i(j) = s_i(j')$ are possible. For valuations and strategies without ties, we will say that the agents have *strict preferences*.

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,

Definition 2.1. A mechanism M is *ordinal* if for any strategy profiles s, 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 $M(\mathrm{s}) = M(\mathrm{s}')$. 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 $(a_1, a_2, \ldots, a_n)$ 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.

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. Formally,

Definition 2.2. A mechanism is *anonymous* if for any input strategy profile $(s_1, s_2, \ldots, s_n)$, every agent i and any permutation $\pi : N \to N$ it holds that $J_i(s_1, s_2, \ldots, s_n) = J_{\pi(i)}(s_{\pi(1)}, s_{\pi(2)}, \ldots, s_{\pi(n)})$. Similarly, a mechanism is *neutral* if for

[3]We can actually show that the assumption $\min_{j \in A} u_i(j) = 0$ is not necessary for our bounds to hold, as long as $\max_{j \in A} u_i(j) = 1$ holds.

any strategy profile $(s_1, s_2, \ldots, s_n)$, every item j and any permutation $\sigma : M \to M$ it holds that $J_i(s_1, s_2, \ldots, s_n) = \sigma^{-1}(J_i(s_1 \circ \sigma, s_2 \circ \sigma, \ldots, s_n \circ \sigma))$, i.e., the mechanism is invariant to the indices of the items.

Note that by this definition, in an anonymous mechanism, agents with exactly the same strategies must have the same probabilities of receiving each item.

Equilibria

An *equilibrium* is a strategy profile in which no agent has an incentive to deviate to a different strategy. We consider five standard equilibrium concepts in this paper: pure Nash, mixed Nash, correlated, coarse correlated and Bayes-Nash equilibria. For the first four, the agents have full information about the valuations of the others. In the Bayesian setting, the valuations are drawn from some distributions and agents know their own valuation and the distributions from which the other valuations are drawn from. We formally define the different equilibrium concepts.

Definition 2.3. 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 ith index. Then a strategy profile $\mathbf{q}$ is called a

(1) *pure Nash equilibrium* if

$$\mathbf{q} = \mathbf{s} \;\text{ and }\; u_i(M_i(\mathbf{s})) \geq u_i(M_i(s_i', \mathbf{s_{-i}})),$$

(2) *mixed Nash equilibrium* if

$$\mathbf{q} = \times_i q_i, \mathbb{E}_{\mathbf{s}\sim\mathbf{q}}[u_i(M_i(\mathbf{s}))] \geq \mathbb{E}_{\mathbf{s_{-i}}\sim q_{-i}}[u_i(M_i((s_i', \mathbf{s_{-i}})))],$$

(3) *correlated equilibrium* if

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

(4) *coarse correlated equilibrium* if

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

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

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

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.

As a relaxation of pure Nash equilibria, we will also consider ϵ-aproximate 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. Finally, a pure strategy profile s is a *dominant strategy equilibrium* if for any agent i, any strategy s_i' and any tuple of strategies $\hat{\mathbf{s}}_{-\mathbf{i}}$ of the other agents, it holds that $u_i(M_i(s_i, \hat{\mathbf{s}}_{-\mathbf{i}}) \geq u_i(M_i(s_i', \hat{\mathbf{s}}_{-\mathbf{i}}))$, i.e. each agent i (weakly) maximizes her utility by using strategy s_i, *regardless* of what the other agents do. A mechanism for which the truth-telling strategy profile $\mathbf{u} = (u_1, \ldots, u_n)$ is always a dominant strategy equilibrium is called *truthful*.

Efficiency

Let $S_{\mathbf{u}}^M$ denote the set of all pure Nash equilibria of mechanism M under true valuation profile u. The measure of efficiency that we will use is the *Price of Anarchy*. The following definition is for the *pure* Price of Anarchy; the definitions with respect to

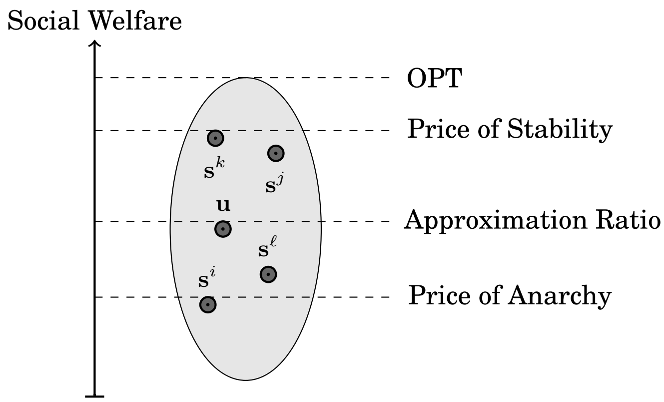

Fig. 1. A pictorial representation of the three different notions of inefficiency. The grey ellipse is the set of all strategy profiles for a valuation profile $\mathbf{u}$ and the profiles $\mathbf{s}^i, \mathbf{s}^j, \mathbf{s}^k$ and $\mathbf{s}^\ell$ are the pure Nash equilibria. Note that if the mechanism in question is truthful, then $\mathbf{u}$ is an equilibrium as well. OPT is the welfare maximizing allocation, which for the purpose of this picture is assumed to be implementable by some (non-equilibrium) strategy (this is true for all the well-known mechanisms). For the inefficiency notions of (a) the Price of Anarchy, (b) the Price of Stability and (c) the Approximation Ratio, we consider (a) the worst equilibrium, (b) the best equilibrium and (c) the truth-telling equilibrium respectively.

other equilibria are analogous.

$$PoA(M) = \sup_{\mathbf{u}\in V^n} \frac{SW_{OPT}(\mathbf{u})}{\min_{\mathbf{s}\in S^M_{\mathbf{u}}} SW_M(\mathbf{u},\mathbf{s})} \tag{1}$$

where $SW_M(\mathbf{u},\mathbf{s}) = \sum_{i=1}^n \mathbb{E}[u_i(M_i(\mathbf{s}))]$ is the expected[4] *social welfare* of mechanism M on strategy profile $\mathbf{s}$ under true valuation profile $\mathbf{u}$, and $SW_{OPT}(\mathbf{u}) = \max_{\mu\in O}\sum_{i=1}^n u_i(\mu_i)$ 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 ith coordinate.

Similarly, we can define a more optimistic measure of inefficiency, the *pure Price of Stability*, which is based on the best equilibrium instead of the worst

$$PoS(M) = \sup_{\mathbf{u}\in V^n} \frac{SW_{OPT}(\mathbf{u})}{\max_{\mathbf{s}\in S^M_{\mathbf{u}}} SW_M(\mathbf{u},\mathbf{s})} \tag{2}$$

The coarse correlated and the Bayesian Price of Anarchy (and Price of Stability) are defined similarly to the pure Price of Anarchy. It is well known that for the first four classes each is contained in the next class, i.e., pure $\subset$ mixed $\subset$ correlated $\subset$ coarse correlated. If we regard the full information setting as a special case of Bayesian setting, we also have pure $\subset$ mixed $\subset$ Bayesian. The hierarchy above implies that for the complete information setting, when proving efficiency guarantees, it suffices to consider the coarse correlated equilibria of a mechanism and in the incomplete information setting, we only need to consider Bayes-Nash equilibria.

If we restrict our attention to the subset of $S^M_{\mathbf{u}}$ that (possibly) contains only the truthtelling equilibrium[5], we obtain the following ratio:

$$ar(M) = \sup_{\mathbf{u}\in V^n} \frac{SW_{OPT}(\mathbf{u})}{SW_M(\mathbf{u})} \tag{3}$$

[4]We remark here that the expectation is with respect to the random choices of the mechanism.

[5]Note that given a valuation profile $\mathbf{u}$, this subset will be either empty, if truthtelling is not an equilibrium, or a singleton, containing $\mathbf{u}$.

where $SW_M(\mathbf{u})$ is a shorthand for $SW_M(\mathbf{u}, \mathbf{u})$. Note that for truthful mechanisms, the set of truth-telling equilibria is always non-empty and the quantity above is called the *approximation ratio* [Procaccia and Tennenholtz 2009]. It should be clear from the definitions that a Price of Anarchy guarantee for a truthful mechanism is stronger than an approximation ratio upper bound, since it bounds the inefficiency of all equilibria, not just the truth-telling ones. On the other hand, a lower bound on the approximation ratio stands "between" a Price of Anarchy and a Price of Stability inapproximability result, in terms of increasing strength (see Figure 1 for a pictorial representation).

3. PRICE OF ANARCHY GUARANTEES

In this section, we will prove the Price of Anarchy guarantees of Probabilistic Serial and Random Priority, for all equilibrium notions up to the coarse correlated equilibria, as well as for the case of incomplete information and the Bayes-Nash equilibria. The results of this section can be summarized as follows: both Probabilistic Serial and Random Priority achieve a Price of Anarchy of $O(\sqrt{n})$ with respect to all the equilibrium notions we consider.

3.1. 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 $O(\sqrt{n})$ for both representations, unit-range and unit-sum. Aziz et al. [2015a] proved that PS has pure Nash equilibria, and so, for ease of exposition, we will first prove the bound for the pure Price of Anarchy and then afterwards, we will explain how to extend the result to more general equilibrium concepts.

We start with the following two lemmas (which hold independently of the choice of representation), which prove that in a pure Nash equilibrium of the mechanism an agent's utility cannot be much smaller 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 $t_j(\mathrm{s})$ be the time when item j is entirely consumed on profile s under $PS(\mathrm{s})$.

LEMMA 3.1. *Let* 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 $t_j(s_i^*, \mathrm{s}_{-\mathbf{i}}) \geq \frac{1}{4} \cdot t_j(\mathrm{s})$.*

PROOF. For ease of notation, let $\mathrm{s}^* = (s_i^*, \mathrm{s}_{-\mathbf{i}})$. Obviously, if $j \succ_{s_i} \ell$ for all $\ell \neq j$ and since all other agents' reports are fixed, $t_j(\mathrm{s}^*) = t_j(\mathrm{s})$ 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 $t_j(\mathrm{s}^*) = 1$ 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 s^*. Obviously, $t_j(\mathrm{s}^*) > \tau$, therefore if $\tau \geq \frac{1}{4} \cdot t_j(\mathrm{s})$ then $t_j(\mathrm{s}^*) \geq \frac{1}{4} \cdot t_j(\mathrm{s})$ and we are done. So assume that $\tau < \frac{1}{4} \cdot t_j(\mathrm{s})$. Next observe that in the interval $[\tau, t_j(\mathrm{s}^*)]$, 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 $\frac{1}{2}+\frac{1}{4}t_j(\mathrm{s})$ so at least $\frac{1}{2}-\frac{1}{4}t_j(\mathrm{s})$ of the item will be consumed by the rest of the agents.

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

$$t_j(\mathrm{s}^*) \ge \beta t_j(\mathrm{s}) \ge t_j(\mathrm{s})\left(\frac{1}{2}-\frac{1}{4}\cdot t_j(\mathrm{s})\right)\frac{1}{\alpha} \ge t_j(\mathrm{s})\left(\frac{1}{2}-\frac{1}{4}\cdot t_j(\mathrm{s})\right) \ge \frac{1}{4}\cdot t_j(\mathrm{s})$$

□

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

LEMMA 3.2. *Let* $\mathbf{u}$ *be the profile of true agent valuations and let* s *be a pure Nash equilibrium. For any agent* i *and any item* j *it holds that the utility of agent* i *at* s *is at least* $\frac{1}{4}\cdot t_j(\mathrm{s})\cdot u_{ij}$.

PROOF. Let $\mathrm{s}' = (s'_i, \mathrm{s}_{-\mathbf{i}})$ be the strategy profile obtained from 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 s is a pure Nash equilibrium, it holds that $u_i(PS_i(\mathrm{s})) \ge u_i(PS_i(\mathrm{s}')) \ge t_j(\mathrm{s}')\cdot u_{ij}$, 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 3.1, it holds that $t_j(\mathrm{s}') \ge \frac{1}{4}\cdot t_j(\mathrm{s})$ and hence $u_i(PS_i(\mathrm{s})) \ge \frac{1}{4}\cdot t_j(\mathrm{s})\cdot u_{ij}$. □

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

THEOREM 3.3. *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 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 $(1/n,\ldots,1/n)$, 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 $p_{ia_1} \ge 1/n$. 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 $p_{ia_1} + p_{ia_2} \ge 2/n$. By a similar argument, for any k, it holds that $\sum_{j=1}^{n} p_{ia_j} \ge k/n$. This implies that regardless of other agents' strategies, agent i can achieve a utility of at least $(1/n)\sum_{j=1}^{n} u_{ij}$. Since s is a pure Nash equilibrium, it holds that $u_i(PS_i(\mathrm{s})) \ge (1/n)\sum_{j=1}^{n} u_{ij}$ as well. Summing over all agents, we get that $SW_{PS}(\mathbf{u},\mathrm{s}) \ge (1/n)\sum_{i=1}^{n}\sum_{j=1}^{n} u_{ij} \ge 1$, which holds for both representations (for unit-sum it holds as an equality). If $SW_{OPT}(\mathbf{u}) \le \sqrt{n}$ then we are done, so assume $SW_{OPT}(\mathbf{u}) > \sqrt{n}$.

Because PS is neutral we can assume $t_j(\mathrm{s}) \le t_{j'}(\mathrm{s})$ for $j < j'$ without loss of generality. Observe that for all $j = 1,\ldots,n$, it holds that $t_j(\mathrm{s}) \ge j/n$. 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 jth item that is entirely consumed, by time $t_j(\mathrm{s})$, 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 3.2, 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 $j = 1, \ldots, k$ 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 $O(\sqrt{n})$. □

Theorem 3.3 establishes the pure Price of Anarchy of Probabilistic Serial. We will next describe how to obtain the same guarantee for all the solution concepts that we consider. Recall that by the equilibrium concept hierarchy mentioned in Section 2, it suffices to extend the theorem to coarse correlated equilibria and Bayes-Nash equilibria.

First, we extend Theorem 3.3 to the case where the solution concept is the coarse correlated equilibrium.

THEOREM 3.4. *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 3.1, 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

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

where the last inequality holds by Lemma 3.1. Using this, we can use very similar arguments to the arguments of the proof of Theorem 3.3 and obtain the bound. □

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

PROOF. The proof is again similar to the proof of Theorem 3.3. Let $\mathbf{u}$ be a valuation profile drawn from some distribution. 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 3.3, 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

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

where the last inequality holds by Lemma 3.1 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

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

and the bound follows. □

3.1.1. Connection to smoothness. Before we conclude the section, we briefly discuss the connection of those extensions with the *smoothness* framework of Roughgarden [2009] (see also [Syrgkanis and Tardos 2013]). According to the definition in [Roughgarden 2009], a game is (λ,μ)-*smooth* if it satisfies the following condition

$$\sum_{i=1}^{n}u_i(s_i^*,\mathbf{s_{-i}}) \geq \lambda SW(\mathbf{s}^*) - \mu SW(\mathbf{s}), \tag{4}$$

where s^* is a pure strategy profile that corresponds to the optimal allocation and 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 3.3 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 different 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 [Roughgarden 2009], 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 [Roughgarden 2009] uses Inequality 4 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 s, when 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 3.3, 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 4. 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.

3.2. Random Priority

Next, we consider the other very well-known mechanism, Random Priority, often referred to as Random Serial Dictatorship, which we will refer to as RP for short. 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.

Before we proceed, we would like to point out that while RP is truthful, it might have other equilibria as well. For example, consider a valuation profile with three agents $1, 2, 3$ and three items a, b, c such that $a \succ_1 b \succ_1 c$, $b \succ_2 c \succ_2 a$ and $b \succ_3 c \succ_3 a$. Clearly, if agents 2 and 3 are being truthful, agent 1 will receive her most-preferred item with certainty a, regardless if she reports $a \succ_1 b \succ_1 c$ truthfully or if she misreports $a \succ_1 c \succ_1 b$ instead and both profiles are pure Nash equilibria. The fact that agent 1 receives item a with certainty is not critical for the example, but in fact, in all the different equilibria, the allocation of the agent will not change (assuming strict preferences) or will change in a way that does not affect the social welfare (assuming weak preferences).

We will prove that the Price of Anarchy of RP is $O(\sqrt{n})$ for both representations and for all the equilibrium concepts we consider. The proof will proceed as follows. First, we will state a general lemma (Lemma 3.7) regarding the approximation ratio of ordinal mechanisms for the unit-range representation and using this lemma, we will then prove that the approximation ratio of RP for unit-range is bounded by $O(\sqrt{n})$ (Lemma 3.8). Next, we will establish a construction that transforms any unit-sum valuation profile to a unit-range valuation profile, maintaining the same asymptotic approximation ratio guarantee, thus bounding the approximation ratio of RP by $O(\sqrt{n})$ for the unit-sum representation as well (Lemma 3.9), which establishes the approximation ratio bound for both representations (Lemma 3.10). Finally, using the observation of the previous paragraph, we will extend the result from the truth-telling equilibria to the set of all equilibria of the mechanism, thus proving the Price of Anarchy guarantee, starting from mixed Nash equilibria (Theorem 3.11) and then extending to coarse-correlated equilibria and Bayes-Nash equilibria (Theorem 3.12).

Remark 3.6. In the following, when referring to the approximation ratio and since we will consider the truth-telling equilibrium, we will denote the input to the mechanism by $\mathbf{u}$ instead of s. Furthermore, for technical reasons, we will consider only valuation vectors when agents have strict preferences. To extend the approximation ratio result of RP to all valuations, the mechanism clearly must be equipped with some tie-breaking rule to settle cases where indifferences appear. For all natural (fixed before the execution of the mechanism) tie-breaking rules the lower bounds still hold. To see this, consider any valuation profile with ties and a tie-breaking rule for random priority. We can add sufficiently small quantities ϵ_{ij} to the valuation profile according to the tie-breaking rule and create a new profile without ties. The assignment probabilities of random priority will be exactly the same as for the version with ties, and random priority achieves the guaranteed approximation ratio on the new profile. Then since ϵ_{ij} were sufficiently small, the same bound holds for the original valuation profile.

3.2.1. Quasi-combinatorial valuation profiles. It will be useful to consider a special class of the unit-range canonically represented valuation functions C_ϵ that we will refer to as *quasi-combinatorial valuation functions*, a straightforward adaptation of a similar notion in [Filos-Ratsikas and Miltersen 2014]. Informally, a valuation function is quasi-combinatorial if the valuations of each agent for every item are close to 1 or close to 0 (the proximity depends on ϵ). Formally,

$$C_\epsilon = \{u \in V | u(A) \subset [0, \epsilon) \cup (1-\epsilon, 1]\},$$

where $u(A)$ is the image of the valuation function u. Let $C_\epsilon^n \subseteq V^n$ be the set of all valuation profiles with n agents whose valuation functions are in C_ϵ. The following lemma implies that when we are trying to prove a lower bound on the approximation ratio of random priority, it suffices to restrict our attention to quasi-combinatorial valuation profiles $C_\epsilon^n \subseteq V^n$ for any value of ϵ.

LEMMA 3.7. *Let M be a truthful, ordinal, anonymous and neutral randomized mechanism for the unit-range representation, and let $\epsilon > 0$. Then*

$$ar(M) = \sup_{\mathbf{u} \in C_\epsilon^n} \frac{SW_{OPT}(\mathbf{u})}{SW_M(\mathbf{u})}.$$

PROOF. Recall that $SW_M(\mathbf{u}) = \mathbb{E}[\sum_{i=1}^n u_i(M_i(\mathbf{u}))]$. Since J is anonymous and neutral, we can assume that the optimal matching is μ^*, where μ^* is the matching with $\mu_i^* = i$ for every agent $i \in N$. Given this, then for any valuation profile $\mathbf{u}$, define

$$g(\mathbf{u}) = \frac{\mathbb{E}[\sum_{i=1}^n u_i(M_i(\mathbf{u}))]}{\sum_{i=1}^n u_i(\mu_i^*)}.$$

Because of this, the approximation ratio can be written as $ar(J) = \sup_{\mathbf{u} \in V^n} \frac{1}{g(\mathbf{u})}$. Now since $C_\epsilon^n \subseteq V^n$, the lemma follows from the following claim:

For all $\mathbf{u} \in V^n$ there exists $\mathbf{u}' \in C_\epsilon^n$ such that $g(\mathbf{u}') \leq g(\mathbf{u})$

We will prove the claim by induction on $\sum_{i=1}^n \#\{u_i(A) \cap [\epsilon, 1-\epsilon]\}$.

Induction basis: Since $\sum_{i=1}^{n} \#\{u_i(A) \cap [\epsilon, 1-\epsilon]\} = 0$, one can clearly see that $u_i \in C_\epsilon$ for all $i \in N$. So, for this case, let $\mathbf{u}' = \mathbf{u}$.

Induction step: Consider a profile $\mathbf{u} \in V^n$ with $\sum_{i=1}^{n} \#\{u_i(A) \cap [\epsilon, 1-\epsilon]\} > 0$. Clearly, there exists an i such that $\#\{u_i(A) \cap [\epsilon, 1-\epsilon]\} > 0$. By this fact, there exist $l, r \in [\epsilon, 1-\epsilon]$, such that $l \le r$, $u_i(A) \subset [0, \epsilon) \cup [l, r] \cup (1-\epsilon, 1]$ and $\{l, r\} \subseteq u_i(A)$.

Let $\bar{l}$ be the largest number such that $\bar{l} \in [0, \epsilon)$ and $\bar{l} \in u_i(A)$. Similarly, let $\bar{r}$ be the smallest number such that $\bar{r} \in (1-\epsilon, 1]$ and $\bar{r} \in u_i(A)$. Note that both those numbers exist, since $\{0, 1\} \subseteq u_i(M)$. Now let $\tilde{l} = \frac{\bar{l}+\epsilon}{2}$, and $\tilde{r} = \frac{\bar{r}+1-\epsilon}{2}$

Now, for any $x \in [\tilde{l} - l, \tilde{r} - r]$, define a valuation function $u_i^x \in V$ as follows:

$$u_i^x(j) = \begin{cases} u_i(j), & \text{for } j \notin u_i^{-1}([\epsilon, 1-\epsilon]) \\ u_i(j) + x, & \text{for } j \in u_i^{-1}([\epsilon, 1-\epsilon]) . \end{cases}$$

This is still a valid valuation function, since by the choice of the interval $[\tilde{l} - l, \tilde{r} - r]$, there are no ties in the image of the function. Let $(u_i^x, \mathbf{u}_{-i})$ be the valuation profile where all agents have the same valuation functions as in $\mathbf{u}$ except for agent i, who has valuation function u_i^x. Define the following function $f : x \to g((u_i^x, \mathbf{u}_{-i}))$. Since J is ordinal, by the definition of function g, we can see that f on the domain $[\tilde{l} - l, \tilde{r} - r]$ is a fractional linear function $x \to (ax + b)/(cx + d)$ for some $a, b, c, d, \in \mathbb{R}$. Since f is defined on the whole interval $[\tilde{l} - l, \tilde{r} - r]$, it is either monotonically increasing, monotonically decreasing or constant in the interval. If f is monotonically increasing, let $\tilde{\mathbf{u}} = (u^{\tilde{l}-l}, \mathbf{u}_{-i})$, otherwise let $\tilde{\mathbf{u}} = (u^{\tilde{r}-r}, \mathbf{u}_{-i})$. Clearly, $g(\tilde{\mathbf{u}}) \le g(\mathbf{u})$ and

$$\sum_{i=1}^{n} \#\{\tilde{u}_i(A) \cap [\epsilon, 1-\epsilon]\} < \sum_{i=1}^{n} \#\{u_i(A) \cap [\epsilon, 1-\epsilon]\}.$$

Then, apply the induction hypothesis on $\tilde{\mathbf{u}}$. This completes the proof. □

The lemma formalizes the intuition that because the mechanism is ordinal, the worst-case approximation ratio is encountered on extreme valuation profiles. Note that truthfulness is only implicitly used in the fact that the truth-telling profile is assumed to be a stable outcome. The lemma also applies to settings where strategic behaviour is not an issue and the loss in welfare is due to other reasons like ordinality or fairness (e.g. see [Aziz et al. 2016; Anshelevich and Sekar 2015; Boutilier et al. 2012]).

The approximation ratio guarantee of RP for the unit-range representation is given by the following lemma.

LEMMA 3.8. *For the unit-range representation, the approximation ratio of RP is $O(\sqrt{n})$.*

PROOF. Because of Lemma 3.7, for the purpose of computing a lower bound on the approximation ratio of random priority, it is sufficient to only consider quasi-combinatorial valuation profiles. Let $\epsilon \le 1/n^3$. Then, there exists $k \in \mathbb{N}$ such that $|k - w^*(\mathbf{u})| \le 1/n^2$, where $w^*(\mathbf{u}) = SW_{OPT}(\mathbf{u})$ is the social welfare of the maximum weight matching on valuation profile $\mathbf{u}$. Since random priority can trivially achieve an expected welfare of 1 (for any permutation the first agent will be matched to her most preferred item), we can assume that $k \ge \sqrt{n}$, otherwise we are done. Note that the maximum weight matching $\mu^* \in O$ assigns k items to agents with $u_i(\mu_i) \in (1-\epsilon, 1]$. Since random priority is anonymous and neutral, without loss of generality we can assume that these agents are $\{1, \ldots, k\}$ and for every agent $j \in N$, it holds that $\mu_j^* = j$. Thus $u_j(j) \in (1-\epsilon, 1]$ for $j = 1, \ldots, k$ and $u_j(j) \in [0, \epsilon)$ for $j = k+1, \ldots, n$.

Consider any run of random priority; one agent is selected in each round. Let $l \in \{0, \ldots, n-1\}$ be any of the n rounds. We will now define the following sets:

$$U_l = \{j \in \{1, \ldots, n\} | \text{ agent } j \text{ has not been selected prior to round } l\}$$
$$G_l = \{j \in U_l | u_j(j) \in (1-\epsilon, 1] \text{ and item } j \text{ is still unmatched}\}$$
$$B_l = \{j \in U_l | u_j(j) \in [0, \epsilon) \text{ or item } j \text{ has already been matched to some agent}\}$$

These three families of sets should be interpreted as three sets that change over the course of the execution of RP. U_l is the set of agents yet to be matched, which is partitioned into G_l, the set of "good" agents, that guarantee a welfare of almost 1 when picked, and B_l, the set of "bad" agents, that do not guarantee any contribution to the social welfare. For the purpose of calculating a lower bound, we will simply bound the sizes of the sets in these families. Obviously, $G_0 = \{1, \ldots, k\}$ and $B_0 = \{k+1, \ldots, n\}$.

The probability that an agent $i \in G_l$ is picked in round l of random priority is $|G_l|/(|G_l| + |B_l|)$, whereas the probability that an agent $i \in B_l$ is picked is $|B_l|/(|G_l| + |B_l|)$. By the discussion above, we can assume that whenever an agent from G_l is picked, her contribution to the social welfare is at least $1-\epsilon$ whereas the contribution from an agent picked from B_l is less than ϵ. In other words, the expected contribution to the social welfare from round l is at least $|G_l|/(|G_l| + |B_l|) - \epsilon$.

We will now upper bound $|G_l|$ and lower bound $|B_l|$ for each l. Consider round l and sizes $|G_l|$ and $|B_l|$. First suppose that some agent i from G_l is picked and the agent is matched with item j. If $j \neq i$ and agent j is in G_l, then $|G_{l+1}| = |G_l| - 2$ and $|B_{l+1}| = |B_l| + 1$, since agent j no longer has its item from the optimal allocation available and so agent j is in B_{l+1}. On the other hand, if $j = i$ or agent j is in B_l then $|G_{l+1}| = |G_l| - 1$ and $|B_{l+1}| = |B_l|$. In either case, $|G_{l+1}| \geq |G_l| - 2$ and $|B_{l+1}| \leq |B_l| + 1$. Intuitively, the picked agent might take an item from a good agent and turn her into a bad agent.

Now suppose that agent i from B_l is picked and the agent is matched with item j. If agent j is in G_l then $|G_{l+1}| = |G_l| - 1$ and $|B_{l+1}| = |B_l|$, since agent j no longer has her item from the optimal allocation available and so agent j is in B_{l+1}. On the other hand, if agent j is in B_l then $|G_{l+1}| = |G_l|$ and $|B_{l+1}| = |B_l| - 1$. In either case, $|G_{l+1}| \geq |G_l| - 2$ and $|B_{l+1}| \leq |B_l| + 1$.

To sum up, in each round l of random priority, we can assume the size of B_l increases by at most 1 and the size of G_l decreases by at most 2. Given this and that $|G_0| = k$ and $|B_0| = n-k$ and that $|G_l| > 0$ for $l \leq \lfloor k/2 \rfloor$, we get

$$\mathbb{E}\left[\sum_{i=1}^{n} u_i(RP_i(\mathbf{u}))\right] \geq \sum_{l=0}^{n}\left(\frac{|G_l|}{|G_l| + |B_l|} - \epsilon\right) \geq \sum_{l=0}^{\lfloor \frac{k}{2} \rfloor} \frac{k-2l}{n-l} - n\epsilon$$

and the ratio is

$$\begin{aligned}
\frac{w^*(\mathbf{u})}{\mathbb{E}\left[\sum_{i=1}^{n} u_i(RP_i(\mathbf{u}))\right]} &\leq \frac{k + \frac{1}{n^2}}{\sum_{l=0}^{\lfloor \frac{k}{2} \rfloor} \frac{k-2l}{n-l} - n\epsilon} \leq \frac{2k}{\sum_{l=0}^{\lfloor \frac{k}{2} \rfloor} \frac{k-2l}{n-l} - n\epsilon} \\
&= \left(\sum_{l=0}^{\lfloor \frac{k}{2} \rfloor} \frac{1 - \frac{2l}{k}}{2(n-l)} - \frac{n\epsilon}{2k}\right)^{-1} \leq \left(\sum_{l=0}^{\lfloor \frac{k}{2} \rfloor} \frac{1 - \frac{2l}{k}}{2n} - \frac{n\epsilon}{2k}\right)^{-1} \\
&\leq \left(\frac{k-11}{8n} - \frac{n\epsilon}{2k}\right)^{-1}.
\end{aligned}$$

The bound is clearly maximized when k is minimized, that is, $k = \sqrt{n}$. Since this bound holds for any $\mathbf{u} \in C_\epsilon^n$, we get

$$ar(RP) = \sup_{\mathbf{u} \in C_\epsilon^n} \frac{w^*(\mathbf{u})}{\mathbb{E}[\sum_{i=1}^n u_i(M_i(\mathbf{u}))]} \leq \left(\frac{\sqrt{n} - 11}{8n} - \frac{n\epsilon}{2\sqrt{n}} \right)^{-1} = \frac{8n}{\sqrt{n} - 11 - 4n\sqrt{n}\epsilon}.$$

We can choose ϵ so that the approximation ratio is at most $20\sqrt{n}$ for $n \geq 400$ and for $n \leq 400$, the bound holds trivially since random priority matches at least one agent with its most preferred item. Overall, the approximation ratio of RP is bounded by $O(\sqrt{n})$. $\square$

Next, we will describe a construction that transforms any unit-sum valuation profile into a unit-range valuation profile, while preserving the $O(\sqrt{n})$ approximation guarantee for RP. This allows us to prove the following lemma.

LEMMA 3.9. *For the unit-sum representation, the approximation ratio of Random Priority is $O(\sqrt{n})$.*

PROOF. Let $\mathbf{u}$ be any unit-sum valuation profile and let C be the constant in the bound from Lemma 3.8. Suppose first that $SW_{OPT}(\mathbf{u}) < 4\sqrt{n}/C$. We will show that random priority guarantees an expected social welfare of 1, which proves the upper bound for this case. Consider any agent i and notice that in random priority, the probability that the agent is picked by the l'th round is l/n, for any $1 \leq l \leq n$ and hence the probability of the agent getting one of its l most preferred items is at least l/n. Let u_i^l be agent i's valuation for its l'th most preferred item; agent i's expected utility for the first round is then at least u_i^1/n. For the second round, in the worst case, agent i's most preferred item has already been matched to a different agent and so the expected utility of the agent for the first two rounds is at least $u_i^1/n + u_i^2/n$. By the same argument, agent i's expected utility after n rounds is at least $\sum_{i=1}^n u_i^l/n = 1/n$. Since this holds for each of the n agents, the expected social welfare is at least 1.

Suppose now that $SW_{OPT}(\mathbf{u}) \geq 4\sqrt{n}/C$. We will transform $\mathbf{u}$ to a unit-range valuation profile $\mathbf{u}''$. First, we will argue that for any valuation profile $\mathbf{u}$, the optimal allocation on $\mathbf{u}$ is a possible outcome of random priority. To see this, first suppose that no agent is matched with her most preferred item in the optimal allocation. Then there must exist agents $i_1, ..., i_k$ such that for each l, agent i_{l+1} is matched with agent i_l's most preferred item and agent i_1 is matched with agent i_k's most preferred item. By swapping items along this cycle, all agents are better off and the allocation is not optimal. Now consider any valuation profile $\mathbf{u}$. Since by the previous argument, there exists an agent j that is matched with her most preferred item j in the optimal allocation for $\mathbf{u}$, random priority could pick this agent first. If we reduce $\mathbf{u}$ by removing the agent i and item j, we obtain a smaller valuation profile $\mathbf{u}'$ where the optimal allocation is the same as in $\mathbf{u}$ but without agent i and item j. Then by inductively applying the same argument, we can verify the claim.[6]

By the argument above, the optimal allocation can be achieved by a run of random priority, so we know that in the optimal allocation at most 1 agent will be matched with her least preferred item. Now consider the valuation profile $\mathbf{u}'$ where each agent i's valuation for her least preferred item is set to 0 (unless it already is 0) and the rest of the valuations are as in $\mathbf{u}$. Since the ordinal preferences of agents are unchanged, random priority performs worse on this valuation profile, and because of the argument of the previous paragraph, it holds that $SW_{OPT}(\mathbf{u}') \geq w^*(\mathbf{u}) - 1/n$. Next consider the

[6] A very similar argument was used in [Bogomolnaia and Moulin 2001] to prove that any Pareto optimal matching is the outcome of a serial allocation.

valuation profile

$$\mathbf{u}'' = \begin{pmatrix} \mathbf{u}' & \mathbf{1} \\ \mathbf{o}^T & 1 \end{pmatrix}$$

where $\mathbf{o} \in \mathbb{R}^n$ and $\mathbf{o}_j = (j-1)/n^5$. That is, $\mathbf{u}''$ has $n+1$ agents and items, where agents $1, ..., n$ have the same valuations for items $1, ..., n$ as in $\mathbf{u}'$, every agent has a valuation of 1 for item $n+1$, and agent $n+1$ only has a significant valuation for item $n+1$. Notice that $\mathbf{u}''$ is a unit-range valuation profile, and $w^*(\mathbf{u}'') \geq w^*(\mathbf{u}')+1$. Furthermore, $\mathbb{E}\left[\sum_{i=1}^n u_i(RP_i(\mathbf{u}'))\right] \geq \mathbb{E}\left[\sum_{i=1}^n u_i(RP_i(\mathbf{u}''))\right] - 2$ and hence

$$\begin{aligned}
\frac{SW_{OPT}(\mathbf{u})}{\mathbb{E}\left[\sum_{i=1}^n u_i(RP_i(\mathbf{u}))\right]} &\leq \frac{SW_{OPT}(\mathbf{u}') + 1/n}{\mathbb{E}\left[\sum_{i=1}^n u_i(RP_i(\mathbf{u}'))\right]} \leq \frac{SW_{OPT}(\mathbf{u}'') + 1/n - 1}{\mathbb{E}\left[\sum_{i=1}^n u_i(RP_i(\mathbf{u}''))\right] - 2} \\
&\leq \left(\frac{\mathbb{E}\left[\sum_{i=1}^n u_i(RP_i(\mathbf{u}''))\right]}{w^*(\mathbf{u}'')} - \frac{2}{w^*(\mathbf{u}'')} \right)^{-1} \leq \left(\frac{C}{\sqrt{n}} - \frac{2}{w^*(\mathbf{u})} \right)^{-1} \\
&\geq \left(\frac{C}{\sqrt{n}} - \frac{2}{4\sqrt{n}/C} \right)^{-1} = \frac{2\sqrt{n}}{C}.
\end{aligned}$$

This completes the proof. □

From Lemma 3.8 and Lemma 3.9, together with the discussion in Remark 3.6, we obtain the following bound on the approximation ratio of RP, which holds regardless of the choice of representation.

LEMMA 3.10. *The approximation ratio of Random Priority is $O(\sqrt{n})$.*

Lemma 3.10 bounds the inefficiency of Random Priority with respect to the truth-telling equilibria. In the following, we will extend the result to the set of all mixed Nash equilibria of the mechanism. Note that the following Theorems are independent of the choice of representation.

THEOREM 3.11. *The mixed Price of Anarchy of Random Priority is $O(\sqrt{n})$.*

PROOF. First, we will prove that if the valuations are distinct, i.e. the preferences are strict, the social welfare is the same in all mixed Nash equilibria of Random Priority. Let i be an agent, and let B be a subset of the items. Let $\mathbf{q}$ 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.

Now, to prove the theorem for the set of all (not necessarily distinct) valuation vectors we proceed as follows. We know from Lemma 3.10 that the social welfare of RP 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 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 s to receive either item. So we can change 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 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. This completes the proof. □

Finally, we extend Theorem 3.11 to the coarse-correlated Price of Anarchy and the Bayesian Price of Anarchy.

THEOREM 3.12. *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 3.11. 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 truthtelling, violating the equilibrium condition. The remaining steps are exactly the same as in the proof of Theorem 3.11.

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 $\mathbb{E}_{\mathbf{u}}\left[\mathbb{E}_{\mathbf{s}\sim\mathbf{q}(\mathbf{u})}\left[u_i(\mathbf{s})\right]\right]$. Using the same argument as the one in the proof of Theorem 3.11, we can lower bound the quantity $\mathbb{E}_{\mathbf{s}\sim\mathbf{q}(\mathbf{u})}\left[u_i(\mathbf{s})\right]$ by $\Omega\left(\frac{SW_{OPT}(\mathbf{u})}{\sqrt{n}}\right)$ and the bound follows. □

4. LOWER BOUNDS

In the previous section, we established the performance guarantees of Probabilistic Serial and Random Priority. In the current section, we will prove lower bounds on the Price of Anarchy of any mechanism, which will allow us to evaluate the quality of those very well-known mechanisms. Since we are interested in mechanisms with good properties, it is natural to consider those mechanisms that have pure Nash equilibria. Unlike the previous section however, the nature of our bounds will depend on the representation.

Interestingly, for the *unit-sum representation*, we will prove a general lower bound of $\Omega(\sqrt{n})$ on the Price of Anarchy of any mechanism, including randomized and cardinal mechanisms. As a corollary, we obtain that both RP and PS are optimal among all mechanisms for the problem. For truthful mechanisms, we establish a stronger result, a similar lower bound with respect to their approximation ratio.

For the *unit-range representation*, we will obtain two different bounds. First, we will prove a lower bound of $\Omega(\sqrt{n})$ on the approximation ratio (and therefore on the Price of Anarchy) of all truthful mechanisms; the bound shows that RP is optimal among all truthful mechanisms for the problem, including cardinal ones. Secondly, we will prove a lower bound of $\Omega(n^{3/4})$ on the Price of Anarchy of all mechanisms (under no restrictions), when the solution concept is the ϵ-equilibrium, for any $\epsilon > 0$.

At the end of the section, we will bound the performance of deterministic mechanisms for both representations, showing that randomization is need for non-trivial approximation guarantees to be achievable.

We start with our general lower bound for unit-sum.

THEOREM 4.1. *For the unit-sum representation, 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 $j \in \{1, \ldots, \sqrt{n}\}$ 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 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 $j = 1, \ldots, \sqrt{n}$, it holds that $p_{i_j j}^{M,\mathbf{s}} \leq 1/\sqrt{n}$ and let $I = \{i_1, i_2, \ldots, i_{\sqrt{n}}\}$. 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 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 $p_{ij}^{M,(s'_i,\mathbf{s}_{-\mathrm{i}})} > p_{ij}^{M,\mathbf{s}}$. 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 s is a pure Nash equilibrium.

Now consider the expected social welfare of M under valuation profile $\mathbf{u}'$ at the pure Nash equilibrium 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. $\square$

We now move on to unit-range and we focus on truthful mechanisms. First we will state a lemma that will be useful for the proof (which holds independently of the representation). These kinds of lemmas are standard in literature (e.g. see [Guo and Conitzer 2010; Filos-Ratsikas and Miltersen 2014]). The lemma implies that when trying to prove lower bounds on the approximation ratio of mechanisms, it suffices to consider mechanisms that are anonymous.

LEMMA 4.2. *For any truthful mechanism M, there exists a truthful, anonymous mechanism M' such that $ar(M') \leq ar(M)$.*

PROOF. Let M' be the mechanism that given any strategy profile $\mathbf{u}$ applies a uniformly random permutation to the set of agents and then applies M on $\mathbf{u}$. The mechanism is clearly anonymous. Furthermore, since $\mathbf{u}$ is a valid input to M, the approximation ratio of M' can not be worse than that of M, since the approximation ratio is calculated over all possible valuation profiles. Since M is truthful and since the permutation is independent of the reports, M' is truthful as well. $\square$

Our bound on the approximation ratio of any truthful mechanism is given by the following lemma.

LEMMA 4.3. *Let M be any truthful mechanism for the unit-range representation. Then, the approximation ratio of M is $\Omega(\sqrt{n})$.*

PROOF. By Lemma 4.2, we can assume that Mechanism M is anonymous. Let $k \geq 2$ be a parameter to be chosen later and let $\mathbf{u} = (u_1, u_2, \ldots, u_n)$ be the unit-range valuation profile where

$$u_i(j) = \begin{cases} 1, & \text{for } j = i \\ \frac{2}{k} - \frac{j}{n}, & \text{for } 1 \leq j \leq k+1, j \neq i \\ \frac{n-j}{n^2}, & \text{otherwise} \end{cases} \qquad \forall i \in \{1, \ldots, k+1\}$$

$$u_i(j) = \begin{cases} 1, & \text{for } j = 1 \\ \frac{2}{k} - \frac{j}{n}, & \text{for } 2 \leq j \leq k+1 \\ \frac{n-j}{n^2}, & \text{otherwise} \end{cases} \qquad \forall i \in \{k+2, \ldots, n\}$$

For $i = 2, \ldots, k+1$, let $\mathbf{u^i} = (u'_i, \mathbf{u_{-i}})$ be the valuation profile where all agents besides agent i have the same valuations as in $\mathbf{u}$ and $u'_i = u_{k+2}$. Note that when agent i on valuation profile $\mathbf{u^i}$, reports u_i instead of u'_i, the resulting valuation profile is $\mathbf{u}$. Since J is anonymous and $u'_i = u_1 = u_{k+2} = \ldots = u_n$, then agent i receives at most a uniform lottery among these agents on valuation profile $\mathbf{u^i}$ and so it holds that

$$\begin{aligned} \mathbb{E}[u'_i(M_i(\mathbf{u^i}))] \;&\leq\; \frac{1}{n-k+1} + \sum_{j=2}^{k+1} \frac{1}{n-k+1}\left(\frac{2}{k} - \frac{j}{n}\right) + \sum_{j=k+2}^{n} \frac{1}{n-k+1} \cdot \frac{n-j}{n^2} \\ &\leq\; \frac{4}{n-k+1} \end{aligned}$$

Next observe that since M is truthful-in-expectation, agent i should not increase her expected utility by misreporting u_i instead of u'_i on valuation profile $\mathbf{u^i}$, that is,

$$\mathbb{E}[u'_i(M_i(\mathbf{u^i}))] \geq \mathbb{E}[u'_i(M_i(\mathbf{u}))] \tag{5}$$

For all $i = 2, \ldots, k+1$, let p_i be the probability that $M_i(\mathbf{u}) = i$. Then, it holds that

$$\mathbb{E}[u'_i(M_i(\mathbf{u}))] \geq p_i \left(\frac{2}{k} - \frac{i}{n}\right) \geq p_i \left(\frac{2}{k} - \frac{k+1}{n}\right)$$

and by Inequality (5) we get

$$\begin{aligned} & p_i \left(\frac{2}{k} - \frac{k+1}{n}\right) \leq \frac{4}{n-k+1} \\ => \;& p_i \leq \frac{4}{n-k+1} \cdot \frac{kn}{2n - k(k+1)} \leq \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2} \end{aligned}$$

Let $p = \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}$, i.e. for all i, $p_i \leq p$. We will next calculate an upper bound on the expected social welfare achieved by M on valuation profile $\mathbf{u}$.

For item $j = 1$, the contribution to the social welfare is upper bounded by 1. Similarly, for each item $j = k+2, \ldots, n$, its contribution to the social welfare is upper bounded by $1/n$. Overall, the total contribution by item 1 and items $k+2, \ldots, n$ will be upper bounded by 2.

We next consider the contribution to the social welfare from items $j = 2, \ldots, k+1$. Define the random variables

$$X_j = \begin{cases} 1, & \text{if } M_j(\mathbf{u}) = j \\ \frac{2}{k} - \frac{j}{n}, & \text{otherwise} \end{cases}$$

The contribution from items $j = 2, \ldots, k+1$ is then $\sum_{j=2}^{k+1} X_j$ and so we get

$$\mathbb{E}\left[\sum_{j=2}^{k+1} X_j\right] = \sum_{j=2}^{k+1} \mathbb{E}\left[X_j\right] \leq \sum_{j=2}^{k+1} \left(p + \frac{2}{k} - \frac{j}{n}\right) \leq kp + 2$$

Overall, the expected social welfare of mechanism M is at most $4 + pk$ while the social welfare of the optimal matching is $k + 1 + \sum_{i=k+2}^{n} \frac{n-i}{n^2}$ which is at least k. Since $p = \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}$, the approximation ratio of J then is

$$ar(M) \geq \left(\frac{4+pk}{k}\right)^{-1} = \left(\frac{4}{k} + \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}\right)^{-1}$$

Let $k = \lfloor\sqrt{n}\rfloor - 1$ and note that $\sqrt{n} - 2 \leq k \leq \sqrt{n} - 1$. Then,

$$\begin{aligned} ar(M) \ &\geq \ \left(\frac{4}{k} + \frac{4}{n-k} \cdot \frac{kn}{2n-(k+1)^2}\right)^{-1} \geq \left(\frac{4}{\sqrt{n}-2} + \frac{4}{n-\sqrt{n}+1} \cdot \frac{(\sqrt{n}-1)n}{2n-(\sqrt{n})^2}\right)^{-1} \\ &\geq \ \left(\frac{4}{\sqrt{n}-2} + \frac{4}{\sqrt{n}}\right)^{-1} \leq \left(\frac{12}{\sqrt{n}} + \frac{4}{\sqrt{n}}\right)^{-1} = \frac{\sqrt{n}}{16}, \end{aligned}$$

The last inequality holds for $n \geq 9$ and for $n < 9$ the bound holds trivially. This completes the proof. □

Note that Lemma 4.3 has the following immediate corollary.

THEOREM 4.4. *For the unit-range representation, the pure Price of Anarchy of any truthful mechanism is* $\Omega(\sqrt{n})$.

Although Theorem 4.1 bounds the Price of Anarchy of *all* mechanisms, including truthful mechanisms, as we mention in Section 2 and in Figure 1, approximation ratio lower bounds are stronger than Price of Anarchy lower bounds. For that reason, and for the sake of completeness, the next lemma shows how to obtain a similar approximation ratio lower bound as the one of Lemma 4.3 for the unit-sum representation as well.

LEMMA 4.5. *Let M be a truthful-in-expectation mechanism for the unit-sum representation. The approximation ratio of M is* $\Omega(\sqrt{n})$.

PROOF. Intuitively, the lemma is true because the valuation profile used in the proof of Lemma 4.3 can be easily modified in a way such that all rows of the matrices of valuations sum up to one. Specifically, consider the following valuation profile:

$$u_i(j) = \begin{cases} 1 - \sum_{j \neq i} u_i(j), & \text{for } j = i \\ \frac{2}{10k} - \frac{j}{10n}, & \text{for } 1 \leq j \leq k+1, j \neq i \\ \frac{n-j}{10n^2}, & \text{otherwise} \end{cases} \qquad \forall i \in \{1, \ldots, k+1\}$$

$$u_i(j) = \begin{cases} 1 - \sum_{j \neq 1} u_i(j), & \text{for } j = 1 \\ \frac{2}{10k} - \frac{j}{10n}, & \text{for } 1 < j \leq k+1 \\ \frac{n-j}{10n^2}, & \text{otherwise} \end{cases} \qquad \forall i \in \{k+2, \ldots, n\}$$

Note that this is exactly the same valuation profile used in the proof of Lemma 4.3 where all entries are divided by ten, except those where the valuation is 1, which are now equal to 1 minus the sum of the valuations for the rest of the items. This modification will only carry a factor of $1/10$ through the calculations and hence the proven bound will be the same asymptotically. □

While the Price of Anarchy of truthful mechanisms for unit-range is settled by Theorem 4.4, quite importantly, it is not clear whether the general lower bound on the Price of Anarchy of all mechanisms that we proved in Theorem 4.1 extends to the unit-range representation as well. In fact, we do not know of any bound for the unit-range case that applies to all mechanisms, 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. While the result applies for any positive ϵ, it is weaker than a corresponding result for exact equilibria.

THEOREM 4.6. *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 $j \in \{2, \ldots, k+1\}$ and 0 for all other items,
— for agent $i \notin I$, i has value δ^3 for items $j \in \{2, \ldots, k+1\}$ and 0 for all other items.

Let $\mathbf{u}$ be such a valuation profile and let s be a Nash equilibrium. In the optimal allocation members of I receive items $\{2, \ldots, k+1\}$ 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 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 s were more than δ. That would imply that the social welfare of 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 s) to receive any item in $\{1, \ldots, k+1\}$ 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 $2, \ldots, k+1$. 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 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 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 $\{2, \ldots, k+1\}$. 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 $\{2, \ldots, k+1\}$. 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 $\{2, \ldots, k+1\}$. 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 $\{2, \ldots, k+1\}$. 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 $\{2, \ldots, k+1\}$. However, since the payoff for any item from $\{2, \ldots, k+1\}$ is only δ^2, such a deviation pays less than δ.

Finally, observe that the social welfare of 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 $\{1,\ldots,k+1\}$. To upper bound the expected social welfare, note that item 1 contributes 1 to the social welfare. Items in $\{2,\ldots,k+1\}$ 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. $n > 1/\delta^4$. □

4.1. Deterministic Mechanisms

Interestingly, if we restrict our attention to *deterministic* mechanisms, then we can prove that only trivial pure Price of Anarchy guarantees are achievable. 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 4.7. *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 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 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 s under true valuation profile $\mathbf{u}$, contradicting the fact that s is a pure Nash equilibrium. □

First, we prove the lower bound for unit-sum.

THEOREM 4.8. *For the unit-sum representation, 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 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 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 $i = 2,\ldots,n$, $u_{i,i-1}' = 1-\epsilon_{i,i-1}'$ 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 $i \in \{2,\ldots,n\}$ has valuation close to 1 for item $i-1$ and small valuations for all other items. Futhermore, she prefers items with smaller indices, except for item $i-1$.

We claim that 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 $\{1,\ldots,i\}$, the same deviation would match her with a more preferred item under $\mathbf{u}$ as well, contradicting the fact that 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. □

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 bound is tight.

Now, using Lemma 4.7, we can then prove the following theorem for unit-range.

THEOREM 4.9. *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 s be a pure Nash equilibrium for this profile and assume without loss of generality that $M_i(\mathbf{s}) = i$. By Lemma 4.7, 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. □

Again, it is not hard to see that the mechanism that naively maximizes the sum of the reported valuations has pure Nash equilibria and achieves the above bound. Interestingly, these lower bounds are the first examples of tight bounds where a different choice of representation results in a different Price of Anarchy bound.

5. PRICE OF STABILITY

In this section, we will consider a more optimistic measure of efficiency, the Price of Stability, which, given a valuation profile, bounds the inefficiency of the best equilibrium instead of the worst. The main result of this section will be with respect to the unit-sum representation; we leave the question of whether similar results hold for unit-range as an open problem. The main result is the following: We extend Theorem 4.1 to the Price of Stability of all mechanisms that are *proportional*, which we will define shortly.

Let $a_1 \succ_i a_2 \succ_i \cdots \succ_i a_n$ 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 $k = 1, 2, \cdots, n$. The notation that we will use for this relation is $p_i \succ_i^{sd} q_i$.

Definition 5.1 (*Safe strategy*). Let M be a mechanism. A strategy s_i is a *safe strategy* if for any strategy profile $\mathbf{s_{-i}}$ of the other players, it holds that $M_i(s_i, \mathbf{s_{-i}}) \succ_i^{sd} \left(\frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n}\right)$.

We will say that a mechanism M has a safe strategy if every agent i has a safe strategy s_i in M. A mechanism M is *proportional* if it has truth-telling as a safe strategy.

We now state our lower bound.

THEOREM 5.2. *For the unit-sum representation, the pure Price of Stability of any proportional mechanism is $\Omega(\sqrt{n})$.*

PROOF. We will actually prove a stronger statement, namely that the Price of Stability of any mechanism with a safe strategy (which might be a different strategy than truth-telling) is bounded by $\Omega(\sqrt{n})$. 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 $j = 1, \cdots, k$ and $u_{ij} = 0$ otherwise.
— For all agents $i \notin I$, let $u_{ii} = 1$ and $u_{ij} = 0, j \neq i$.

Now let 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 s is

$$\mathbb{E}[u_i(\mathbf{s})] \; = \sum_{j\in[n]} p^{\mathbf{s}}_{ij} v_{ij} \geq \sum_{j\in[n]} p^{(s'_i,\mathbf{s_{-i}})}_{ij} v_{ij} \geq \frac{1}{n}\sum_{j\in[n]} v_{ij} = \frac{1}{n},$$

due to the fact that 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 $\mathbb{E}[u_i(\mathbf{s})] = \sum_{j\in[n]} p^{\mathbf{s}}_{ij} v_{ij} = (\sum_{j=1}^{k} p^{\mathbf{s}}_{ij})/k$. Because s is a pure Nash equilibrium, it holds that $\mathbb{E}[u_i(\mathbf{s})] \geq 1/n$, so we get that $\sum_{j=1}^{k} p^{\mathbf{s}}_{ij} \geq k/n$ for all $i \in I$. As for the rest of the agents,

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

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 $k = \sqrt{n}$. □

Due to Theorem 5.2, in order to obtain an $\Omega(\sqrt{n})$ bound for a mechanism M, it suffices to prove that M has is proportional. In fact, most reasonable mechanisms, including Random Priority and Probabilistic Serial, as well as all ordinal *envy-free* mechanisms are proportional. We start with a definition.

Definition 5.3 (*Envy-freeness*). A mechanism M is (ex-ante) *envy-free* if for all agents i and r and all profiles s, it holds that $\sum_{j=1}^{n} p^{\mathbf{s}}_{ij} s_{ij} \geq \sum_{j=1}^{n} p^{\mathbf{s}}_{rj} s_{rj}$. Furthermore, if M is ordinal, then this implies $p_i^{M,\mathbf{s}} \succ^{sd}_{s_i} p_r^{M,\mathbf{s}}$.

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

LEMMA 5.4. *Let M be an ordinal, envy-free mechanism. Then M is proportional.*

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 $\sum_{j=1}^{\ell} p^{\mathbf{s}}_{ij} \geq \sum_{j=1}^{\ell} p^{\mathbf{s}}_{rj}$ for all agents $r \in \{1, \ldots, n\}$ and all $\ell \in \{1, \ldots, n\}$. Summing up these inequalities for agents $r = 1, 2, \ldots, n$ we obtain

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

which implies that $\sum_{j=1}^{\ell} p^{\mathbf{s}}_{ij} \geq \frac{\ell}{n}$, for all $i \in \{1, \ldots, n\}$, and for all $\ell \in \{1, \ldots, n\}$. □

Note that since Probabilistic Serial is ordinal and envy-free [Bogomolnaia and Moulin 2001], by Lemma 5.4, it is proportional and hence Theorem 5.2 applies. It is not hard to see that Random Priority is proportional too.

LEMMA 5.5. *Random Priority is proportional.*

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 $\ell = 1, \ldots, n$, it holds that $\sum_{i=1}^{\ell} p_{ij} \geq \ell/n$. □

Note that the safe strategy condition is in a sense a minimal condition required for Theorem 5.2, because we can not hope to prove a strong upper bound on the price of stability of all mechanisms. From the discussion below, it should also be clear that the safe strategy property does not imply truthfulness.

To see why the statement above is true, consider the following deterministic, *Randomly Dictatorial* mechanism RD [Svensson 1999]: Select an agent i^* uniformly at random and match her with her most preferred item j^*. Then, fix an ordering of the rest of the agents and match them sequentially according to this ordering to items

$$j_1 \in \arg\max_{j \in A \setminus \{j^*\}} u_{i^*j} \;\;,\;\; j_2 \in \arg\max_{j \in A \setminus \{j^*, j_1\}} u_{i^*j}$$

and so on, breaking ties arbitrarily. Note that this mechanism is truthful; once agent i^* is selected, she is matched with her most preferred item and the rest of the agents can not influence the outcome. However, it is easy to see that the mechanism has other equilibria as well; any report such that j^* is on top of agent i^*'s preference ranking grants the agent maximum utility. In particular, there is some strategy s_{i^*} of agent i^* that results in an welfare-optimal assignment for the rest of the agents. We can prove the following theorem.[7]

THEOREM 5.6. *For the unit-sum representation, the Price of Stability of the Randomly Dictatorial mechanism RD is at most 2.*

PROOF. Consider any valuation profile $\mathbf{u}$ and assume first that $SW_{OPT}(\mathbf{u}) \geq 2$. Given the choice of some agent i and her most preferred item j, in the best Nash equilibrium, the mechanism outputs a social welfare optimal matching $OPT_{-i}(\mathbf{u_{-i}})$ for agents in $N \backslash \{i\}$ and items in $A \backslash \{j\}$. Since $OPT_{-i}(\mathbf{u_{-i}})$ is optimal, it is at least as good as the matching that matches every agent $l \in N \backslash \{i, i'\}$ with item $OPT_l(\mathbf{u})$, except agent i', the agent for which $OPT_{i'}(\mathbf{u}) = j$, who is matched with item $OPT_i(\mathbf{u})$. In other words, for every realization of randomness, the mechanism produces a matching that is at least as good as $OPT(\mathbf{u})$, except for the allocation of two agents that is swapped: the agent i chosen by the mechanism and the agent that receives agent i's most preferred item in the optimal matching.

Let v_i be the valuation of agent i for her most preferred item and let w_i be her valuation for item $OPT_i(\mathbf{u})$. Then, from the discussion above, it holds that for every choice of agent i (with most preferred item j), the welfare achieved is at least $SW_{OPT}(\mathbf{u}) + v_i - w_i - w_j$, which is at least $SW_{OPT}(\mathbf{u}) - 1$, since $v_i \geq w_i$ and $w_j \leq 1$. The Price of Stability is then at most $SW_{OPT}(\mathbf{u})/(SW_{OPT}(\mathbf{u}) + 1)$ which is at most 2, since $SW_{OPT}(\mathbf{u}) \geq 2$.

Assume from now on that $SW_{OPT}(\mathbf{u}) \leq 2$. Observe that, in the best Nash equilibrium of the Randomly Dictatorial mechanism RD on input $\mathbf{u}$, the outcome is at least as good as the outcome of Random Priority; in particular, there exists some Nash equilibrium such that $RD(\mathbf{u}) = RP(\mathbf{u})$. By the proof of Lemma 3.9, we know that, for the unit-sum representation, $SW_{RP}(\mathbf{u}) \geq 1$ and hence $SW_{RD}(\mathbf{u}) \geq 1$ as well. Since $SW_{OPT}(\mathbf{u}) \leq 2$, this proves the theorem. □

It is not hard to see that the Price of Anarchy of the Randomly Dictatorial mechanism is $\Theta(n)$. Given that bestowing the rights to the allocation to a single agent is intuitively not a good choice, the above result indicates that one should perhaps be careful when

[7]It is not hard to see that a simple deterministic dictatorial mechanism achieves the same Price of Stability guarantee for unit-range as well.

adopting the Price of Stability as the measure of performance. The same idea applies in general to other problems, where “satisfying” a single agent can not hurt the social objective very much; the Price of Stability was suggested as a measure of inefficiency in the absence of a central planner but the existence of such equilibria like the one above essentially delegate the central planning task to a single agent.

6. CONCLUSION AND FUTURE WORK

In this paper, we studied the problem of social welfare maximization in the fundamental setting of one-sided matching and explored the capabilities and limitations of all mechanisms for the problem.

Our results are rather negative: we identify non-constant lower bounds 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. For example, one could attempt 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. Recent experimental studies [Hosseini et al. 2016] attempt to quantify the performance of one-sided matching mechanisms on typical valuation profiles but do not consider the equilibrium behaviour of mechanisms like Probabilistic Serial; this is something worth investigating.

Caragiannis et al. [2016] have recently proposed a resource augmentation framework, introduced on a problem which they refer to as facility assignment, which is essentially an one-sided matching setting on a metric space with costs instead of utilities. The idea is to compare the optimal mechanism with the mechanism in question, when the two mechanism operate under a different set of resources; in our context that would mean that mechanisms like Probabilistic Serial or Random Priority would have access to multiple copies of the items. The framework was proposed for truthful mechanisms, but it can easily be applied to all mechanisms in terms of their equilibrium behaviour. Conversely, our lemmas for bounding the Price of Anarchy of Probabilistic Serial could possibly be applied to bound the inefficiency of the mechanism for the version of problem they consider as well.

In a recent paper, Aziz et al. [2016] study the inefficiency of one-sided matching mechanisms for the objective of *egalitarian social welfare* maximization, i.e. maximizing the utility of the least satisfied agent. They prove bounds on the approximation ratio of matching mechanisms when the constraint is truthfulness, ordinality or envy-freeness, but in the latter two cases, the bounds are proven under the assumption that strategic considerations are not present. There is a clear open question there, analogous to what we do here: “What is the best mechanism in terms of the Price of Anarchy, for egalitarian welfare maximization?”. The same question can be posed for any sensible aggregate objective for one-sided matching, including the ordinal measures of efficiency that we mentioned in the introduction, which are typically studied with respect to the approximation ratio of truthful mechanisms only.

In general, one could adopt the following agenda in mechanism design without money:

— Choose a *problem*, e.g. indivisible item allocation, divisible item allocation, cake cutting, machine scheduling, etc.

— Choose an *objective*, like *cardinal* objectives, e.g. (utilitarian) social welfare, egalitarian social welfare or makespan, or *ordinal* objectives, e.g. Borda scores, matching cardinality or rank-approximation.
— Try to find the *best mechanisms* for each problem in terms of the Price of Anarchy, the Price of Stability or in the case of truthful mechanisms, the Approximation Ratio.

It seems natural to study the efficiency of mechanisms for a problem in a unified framework. For example, in divisible item allocation and utilitarian welfare maximization, there is work bounding the approximation ratio of truthful mechanisms [Han et al. 2011] and the Price of Anarchy of non-truthful mechanisms [Feldman et al. 2009; Brânzei et al. 2014] but the connection between the results is not clearly established. Furthermore, the question whether non-truthful mechanisms for divisible item allocation can outperform truthful ones in equilibrium is yet not answered; a unified approach would first state those questions explicitly and then try to answer them.

Interestingly, contrary to our setting here for which the answer to the corresponding question is "no", recently Giannakopoulos et al. [2016], proved that in machine scheduling and the makespan objective, there exist non-truthful mechanisms which significantly outperform the best truthful ones, whose limitations were established in [Koutsoupias 2014].

Our investigations in this paper shed much light on the question of the efficiency of one-sided matching mechanisms, but there are still some interesting open problems. The main 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. An almost identical question is the following: "Is there a cardinal mechanism that achieves a Price of Anarchy guarantee better than $O(\sqrt{n})$ for unit-range?" From our results, we know that the answer to the corresponding question is "no" for unit-sum, as well as for any truthful mechanism for both representations. If the answer for unit-range turns out to be "yes" that would show that unlike the case of truthful mechanisms, the choice of representation actually matters for the Price of Anarchy bounds. A candidate mechanism for attempting to answer the question above in the positive could be the pseudo-market mechanism of Hylland and Zeckhauser [1979].

ACKNOWLEDGMENTS

The authors would like to thank Piotr Krysta and Peter Bro Miltersen for useful discussion.

REFERENCES

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

Atila Abdulkadiroglu and Tayfun Sönmez. 2013. Matching markets: Theory and practice. *Advances in Economics and Econometrics* 1 (2013), 3–47.

Marek Adamczyk, Piotr Sankowski, and Qiang Zhang. 2014. Efficiency of Truthful and Symmetric Mechanisms in One-sided Matching. In *Algorithmic Game Theory*. Springer, 13–24.

Elliot Anshelevich and Sanmay Das. 2010. Matching, cardinal utility, and social welfare. *ACM SIGECom Exchanges* 9, 1 (2010), 4.

Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. 2008. The price of stability for network design with fair cost allocation. *SIAM J. Comput.* 38, 4 (2008), 1602–1623.

Elliot Anshelevich and Shreyas Sekar. 2015. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. (2015).

Haris Aziz, Jiashu Chen, Aris Filos-Ratsikas, Simon Mackenzie, and Nicholas Mattei. 2016. Egalitarianism of random assignment mechanisms. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Nina Narodytska, and Toby Walsh. 2015a. Equilibria Under the Probabilistic Serial Rule. *arXiv preprint arXiv:1502.04888* (2015).

Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Nina Narodytska, and Toby Walsh. 2015b. Manipulating the probabilistic serial rule. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1451–1459.

Salvador Barbera. 2010. Strategy-proof Social Choice. In *Handbook of Social Choice and Welfare*, K. J. Arrow, A. K. Sen, and K. Suzumura (Eds.). Vol. 2. North-Holland: Amsterdam, Chapter 25.

Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. 2011. The price of fairness. *Operations research* 59, 1 (2011), 17–31.

Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. 2011. Social Welfare in One-Sided Matching Markets without Money. In *APPROX-RANDOM*. 87–98.

Kshipra Bhawalkar and Tim Roughgarden. 2011. Welfare guarantees for combinatorial auctions with item bidding. In *SODA '11: Proceedings of the Twenty-Second annual ACM-SIAM symposium on Discrete algorithms*. SIAM.

Anna Bogomolnaia and Hervé Moulin. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100 (2001), 295–328.

Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. 2012. Optimal social choice functions: a utilitarian view. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 197–214.

Steven J Brams, Michal Feldman, John K Lai, Jamie Morgenstern, and Ariel D Procaccia. 2012. On Maxsum Fair Cake Divisions.. In *AAAI*.

Simina Brânzei, Yiling Chen, Xiaotie Deng, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. 2014. The Fisher market game: Equilibrium and welfare. In *Twenty-Eighth Conference on Artificial Intelligence*. 587–593.

Ioannis Caragiannis, Aris Filos-Ratsikas, Soren Kristoffer Stiil Frederiksen, Kristoffer Arnsfelt Hansen, and Zihan Tan. 2016. Truthful Facility Assignment with Resource Augmentation: An Exact Analysis of Serial Dictatorship. *arXiv preprint arXiv:1602.08023* (2016).

Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2012. The efficiency of fair division. *Theory of Computing Systems* 50, 4 (2012), 589–610.

Ioannis Caragiannis, Swaprava Nath, Ariel D Procaccia, and Nisarg Shah. 2016. Subset Selection Via Implicit Utilitarian Voting. *25th International Joint Conference on Artificial Intelligence (IJCAI-16)* (2016), 151–157.

Ioannis Caragiannis and Ariel D Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175, 9 (2011), 1655–1671.

Deeparnab Chakrabarty and Chaitanya Swamy. 2014. Welfare maximization and truthfulness in mechanism design with ordinal preferences. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*. 105–120.

Yun Kuen Cheung. 2016. Better Strategyproof Mechanisms without Payments or Prior—An Analytic Approach. *arXiv preprint arXiv:1604.05243* (2016).

George Christodoulou, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Paul W Goldberg, Jie Zhang, and Jinshan Zhang. 2016. Social Welfare in One-sided Matching Mechanisms. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

George Christodoulou, Annamária Kovács, and Michael Schapira. 2008. Bayesian Combinatorial Auctions. In *ICALP '08: Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*. Springer-Verlag.

George Christodoulou, Annamária Kovács, Alkmini Sgouritsa, and Bo Tang. 2016. Tight Bounds for the Price of Anarchy of Simultaneous First-Price Auctions. *ACM Transactions on Economics and Computation* 4, 2 (2016), 9.

Yuga J Cohler, John K Lai, David C Parkes, and Ariel D Procaccia. 2011. Optimal Envy-Free Cake Cutting.. In *AAAI*.

Richard Cole, Vasilis Gkatzelis, and Gagan Goel. 2013a. Mechanism design for fair division: allocating divisible items without payments. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 251–268.

Richard Cole, Vasilis Gkatzelis, and Gagan Goel. 2013b. Positive results for mechanism design without money. In *Proceedings of the 12th International conference on Autonomous Agents and Multi-agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1165–1166.

Hervé Crès and Hervé Moulin. 2001. Scheduling with opting out: improving upon random priority. *Operations Research* 49, 4 (2001), 565–577.

Bart de Keijzer, Evangelos Markakis, Guido Schäfer, and Orestis Telelis. 2013. On the Inefficiency of Standard Multi-Unit Auctions. In *ESA'13: Proceedings of the 21st European conference on Algorithms*.

Ozgün Ekici and Onur Kesten. 2010. *On the Ordinal Nash Equilibria of the Probabilistic Serial Mechanism*. Technical Report. working paper, Tepper School of Business, Carnegie Mellon University.

Uriel Feige and Moshe Tennenholtz. 2010. Responsive lotteries. In *Algorithmic Game Theory*. Springer, 150–161.

Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. 2013. Simultaneous Auctions are (almost) Efficient. In *STOC '13: Proceedings of the 45th symposium on Theory of Computing*.

Michal Feldman, Kevin Lai, and Li Zhang. 2009. The proportional-share allocation market for computational resources. *Parallel and Distributed Systems, IEEE Transactions on* 20, 8 (2009), 1075–1088.

Aris Filos-Ratsikas. 2015. *Social Welfare in Algorithmic Mechanism Design Without Money*. Ph.D. Dissertation. Department Office Computer Science, Aarhus University.

Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. 2014. Social welfare in one-sided matchings: Random priority and beyond. In *7th International Symposium on Algorithmic Game Theory (SAGT)*. 1–12.

Aris Filos-Ratsikas and Peter Bro Miltersen. 2014. Truthful approximations to range voting. (2014), 175–188.

Duncan K Foley. 1967. Resource allocation and the public sector. *YALE ECON ESSAYS, VOL 7, NO 1, PP 45-98, SPRING 1967. 7 FIG, 13 REF.* (1967).

David Gale and Lloyd S Shapley. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.

Yiannis Giannakopoulos, Elias Koutsoupias, and Maria Kyropoulou. 2016. The Anarchy of Scheduling Without Money. *9th International Symposium on Algorithmic Game Theory (SAGT)* (2016), to appear.

Mingyu Guo and Vincent Conitzer. 2010. 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*. 881–888.

Li Han, Chunzhi Su, Linpeng Tang, and Hongyang Zhang. 2011. On strategy-proof allocation without payments or priors. In *Proceedings of the 7th International Workshop on Internet and Network Economics*. Springer, 182–193.

Tadashi Hashimoto, Daisuke Hirata, Onur Kesten, Morimitsu Kurino, and M Utku Ünver. 2014. Two axiomatic approaches to the probabilistic serial mechanism. *Theoretical Economics* 9, 1 (2014), 253–277.

Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. 2011. Non-price equilibria in markets of discrete goods. In *EC '11: Proceedings of the 12th ACM conference on Electronic commerce*. ACM.

Daniel M. Hausman. 1995. The impossibility of interpersonal utility comparisons. *Mind* (1995), 473–490.

Hadi Hosseini, Kate Larson, and Robin Cohen. 2015. Matching with Dynamic Ordinal Preferences.. In *AAAI*. 936–943.

Hadi Hosseini, Kate Larson, and Robin Cohen. 2016. Investigating the Characteristics of One-Sided Matching Mechanisms. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1443–1444.

Aanund Hylland and Richard Zeckhauser. 1979. The Efficient Allocation of Individuals to Positions. *The Journal of Political Economy* 87, 2 (1979), 293–314.

Jeremy Karp, Aleksandr M Kazachkov, and Ariel D Procaccia. 2014. Envy-Free Division of Sellable Goods. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.

Akshay-Kumar Katta and Jay Sethuraman. 2006. A solution to the random assignment problem on the full preference domain. *Journal of Economic theory* 131, 1 (2006), 231–250.

Onur Kesten. 2006. *Probabilistic serial and top trading cycles from equal division for the random assignment problem*. Technical Report. mimeo.

Fuhito Kojima and Mihai Manea. 2010. Incentives in the probabilistic serial mechanism. *Journal of Economic Theory* 145, 1 (2010), 106–123.

Elias Koutsoupias. 2014. Scheduling without payments. *Theory of Computing Systems* 54, 3 (2014), 375–387.

Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-case equilibria. In *STACS 99*. Springer, 404–413.

Piotr Krysta, David Manlove, Baharak Rastegari, and Jinshan Zhang. 2014. Size versus truthfulness in the House Allocation problem. In *Proceedings of the 15th ACM Conference on Economics and Computation*. ACM, 453–470.

Timo Mennle and Sven Seuken. 2014. An axiomatic approach to characterizing and relaxing strategyproofness of one-sided matching mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation*. 37–38.

Ariel D Procaccia and Moshe Tennenholtz. 2009. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce*. ACM, 177–186.

Tim Roughgarden. 2009. Intrinsic robustness of the price of anarchy. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*. ACM, 513–522.

Tim Roughgarden. 2014. Barriers to near-optimal equilibria. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. IEEE, 71–80.

Lloyd Shapley and Herbert Scarf. 1974. On cores and indivisibility. *Journal of mathematical economics* 1, 1 (1974), 23–37.

Tayfun Sönmez and M Utku Ünver. 2005. House allocation with existing tenants: an equivalence. *Games and Economic Behavior* 52, 1 (2005), 153–185.

Tayfun Sönmez and M Utku Ünver. 2011. Matching, allocation, and exchange of discrete resources. *Handbook of social Economics* 1 (2011), 781–852.

Lars-Gunnar Svensson. 1999. Strategy-proof allocation of indivisble goods. *Social Choice and Welfare* 16, 4 (1999), 557–567.

Vasilis Syrgkanis and Eva Tardos. 2013. Composable and Efficient Mechanisms. In *STOC '13: Proceedings of the 45th symposium on Theory of Computing*.

UNOS. 1984. United Network for Organ Sharing. http://www.unos.org/. (1984). Online; accessed 27-May-2016.

John Von Neumann and Oskar Morgenstern. 1944. *Theory of games and economic behavior (60th Anniversary Commemorative Edition 2007)*. Princeton university press.

Lin Zhou. 1990. On a Conjecture by Gale about One-Sided Matching Problems. *Journal of Economic Theory* 52 (1990), 123–135.