

Average-case Analysis of the Assignment Problem with Independent Preferences

Yansong Gao¹ and Jie Zhang²

¹Applied Mathematics and Computational Science, University of Pennsylvania

²Electronics and Computer Science, University of Southampton

gaoyans@sas.upenn.edu, jie.zhang@soton.ac.uk

Abstract

The fundamental assignment problem is in search of welfare maximization mechanisms to allocate items to agents when the private preferences over indivisible items are provided by self-interested agents. The mainstream mechanism *Random Priority* is asymptotically the best mechanism for this purpose, when comparing its welfare to the optimal social welfare using the canonical *worst-case approximation ratio*. Despite its popularity, the efficiency loss indicated by the worst-case ratio does not have a constant bound [Filos-Ratsikas *et al.*, 2014]. Recently, [Deng *et al.*, 2017] show that when the agents' preferences are drawn from a uniform distribution, its *average-case approximation ratio* is upper bounded by 3.718. They left it as an open question of whether a constant ratio holds for general scenarios. In this paper, we offer an affirmative answer to this question by showing that the ratio is bounded by $1/\mu$ when the preference values are independent and identically distributed random variables, where μ is the expectation of the value distribution. This upper bound also improves the upper bound of 3.718 in [Deng *et al.*, 2017] for the Uniform distribution. Moreover, under mild conditions, the ratio has a *constant* bound for any independent random values. En route to these results, we develop powerful tools to show the insights that in most instances the efficiency loss is small.

1 Introduction

Studies in Mechanism Design focus on designing mechanisms in which truth-telling is a dominant strategy, and so the rational, risk-neutral agents are motivated to play their truth-telling strategies. Subject to this qualitative constraint, a *truthful mechanism* quantitatively optimizes an objective such as social welfare or revenue. Build upon the classical Mechanism Design framework, Algorithmic Mechanism Design [Nisan and Ronen, 1999] advances research and has evolved to employing two typical analytic tools in Computer Science. One of these tools imposes polynomial-time implementability on designing mechanisms, and the other evaluates mechanism performances through the lens of the worst-

case analysis. Typically, the (*worst-case*) *approximation ratio* assesses to what extent a mechanism approximately optimizes an objective over all possible inputs and it is dominated by the worst-case inputs [Procaccia and Tennenholtz, 2009]. There is now an extensive literature on topics including matching [Dughmi and Ghosh, 2010; Anshelevich, 2016], scheduling [Koutsoupias, 2014; Christodoulou *et al.*, 2009], facility location [Filos-Ratsikas *et al.*, 2015], kidney exchange [Hajaj *et al.*, 2015], fair division [Chen *et al.*, 2013], social choice [Anshelevich *et al.*, 2018], and auction design [Hartline and Roughgarden, 2009; Mu'alem and Nisan, 2008; Archer *et al.*, 2003]. For a more detailed summary, we refer the reader to [Nisan *et al.*, 2007].

In the *assignment problem* (a.k.a., *one-sided matching* or *house allocation* problem), there is a set of agents and a set of items. The agents participate in a mechanism by reporting their private preferences over the items. The mechanism then assigns items to agents, according to a pre-defined allocation function. The folklore mechanism *Random Priority* (a.k.a., *Random Serial Dictatorship*) is mainstream because it satisfies appealing properties including anonymity, truthfulness, and ex-post Pareto efficiency. In addition, there exists no mechanism that is ex-ante Pareto efficient while keeping the first two desired properties [Zhou, 1990]. *Random Priority* and its variants have notable practical applications. For example, they are used in United States Naval Academy placement [Roth and Sotomayor, 1990], social or government-subsidised housing [Abdulkadiroglu and Sönmez, 1999], graduate housing allocation at a large number of universities [Abdulkadiroglu and Sönmez, 1998], and high school student assignment in New York [Pathak and Sethuraman, 2011].

For the purpose of maximizing *social welfare*, i.e., the sum of all agents' utilities, *Random Priority* can only achieve a $\Theta(1/\sqrt{n})$ fraction of the optimal social welfare in the worst case, where n is the number of agents and items, and it is asymptotically the best amongst all truthful mechanisms [Filos-Ratsikas *et al.*, 2014]. This negative result was considered a cautionary tale discouraging the wide applications of *Random Priority*. Fortunately, the smoothed analysis and average-case analysis mitigate it by revealing positive results on social welfare approximation. In [Deng *et al.*, 2017], the authors show that when the worst-case inputs are subject to small random noise, *Random Priority* attains social welfare within a constant factor of the optimal welfare. In addition,

when agents' preferences are drawn from a uniform distribution, on average, the optimal social welfare is no more than $1 + e$ times of the social welfare attainable by Random Priority, where e is the Euler's number. Therefore, on average, the efficiency loss is small. The primal question that remains open is: Does the constant average-case ratio result hold for general probability distributions? In this paper, we partially answer this question by showing the following results.

- When agents' preference values are independent and identically distributed random variables, the average-case approximation ratio is upper bounded by $1/\mu$, where μ is the expectation of these random variables.
- When agents' preferences values are independent but not necessarily identically distributed random variables, under mild conditions, the average-case approximation ratio is upper bounded by a constant.

We note that the $1/\mu$ upper bound for the i.i.d. case improves the 3.718 upper bound for the Uniform distribution $\mathcal{U}[0, 1]$, as its expectation $\mu = \frac{1}{2}$ and $1/\mu$ yields an upper bound of 2. Taken together, these results further pin down the wide-applicability of the Random Priority mechanism, for which the worst-case analysis is insufficient to manifest.

To achieve these results, we employ the Central Limit Theorems and carefully calibrate the values of some parameters in a few building blocks to ascertain the rate that the properly normalized sum of preference values converges to a Normal distribution. By successfully compositing these building blocks together, we bound the average-case ratios.

1.1 Related Work

The one-sided matching problem was introduced in [Hylland and Zeckhauser, 1979] and has been studied extensively ever since. Over the years, several different mechanisms have been proposed with various desirable properties related to truthfulness, fairness and economic efficiency with Probabilistic Serial [Abdulkadiroğlu and Sönmez, 2013; Bhalgat *et al.*, 2011; Bogomolnaia and Moulin, 2001; Katta and Sethuraman, 2006] and Random Priority [Abdulkadiroglu and Sönmez, 1998; Svensson, 1999; Christodoulou *et al.*, 2016; Aziz *et al.*, 2013] being the two prominent examples. In the indivisible goods setting, the Top Trading Cycles (TTC) method is well-studied and generalized to investigate various problems. In particular, [Abdulkadiroglu and Sönmez, 1998] proposed an adaptation of the TTC method and established an equivalence between the adapted mechanism and Random Priority. [Kesten, 2009] proposed several extensions of these popular mechanisms and presented an equivalence result between those mechanisms in terms of economic efficiency. In the presence of incentives, the assignment problem was extensively studied in Computer Science and Economics over the years [Zhou, 1990; Dughmi and Ghosh, 2010; Mennle and Seuken, 2014]. We refer the interested reader to surveys [Abdulkadiroğlu and Sönmez, 2013; Sönmez and Ünver, 2011]. [Bhalgat *et al.*, 2011] studied the approximation ratio of matching mechanisms, when the objective is maximization of ordinal social welfare, a notion of efficiency that they define based solely on ordinal information.

Under the average-case analysis framework, [Zhang, 2018] tackled the scheduling unrelated machines problem and showed that the average-case approximation ratio of the mechanism devised in [Koutsoupias, 2014] is upper bounded by a constant, when the machines' costs follow any independent and identical distribution.

A similar notion but fundamentally different to our approach exists in Bayesian analysis [Chawla and Sivan, 2014; Hartline and Lucier, 2010]. We leave a more detailed discussion of the two notions to the next section after formally present the definition of the average-case approximation ratio, to highlight the comparison.

2 Preliminaries

We study the one-sided matching problem that consists of n agents and n indivisible items. The agents are endowed with von Neumann - Morgenstern utilities over the items. Throughout the paper, we denote the utility derived by agent i on obtaining a unit amount of item j by a_{ij} . In particular, following the classical literature [Zhou, 1990; Barbera, 2010], agents' preferences a_{ij} are represented by *unit-range* values. That is, with normalization, any agent i 's valuation on its most preferred item is 1, i.e., $\max_j \{a_{ij}\} = 1$, and its valuation on the least preferred item is 0, i.e., $\min_j \{a_{ij}\} = 0$. However, note that our model would be more general and some calculations would be cleaner if we drop these constraints but only require that $0 \leq a_{ij} \leq 1, \forall i, j \in [n]$. A *valuation profile* (or interchangeably, an *instance*) of agents' preferences can be represented by a matrix $A = [a_{ij}]_{n \times n}$, where the row vector (a_{i1}, \dots, a_{in}) indicates the valuation of agent i 's preference.

A matching mechanism collects agents' preference valuations and output an assignment of items to them. Denote a matching assignment by a matrix $X = [x_{ij}]_{n \times n}$, where x_{ij} indicates the probability of agent i receiving item j . So, a mechanism is a mapping from the input instance A to an output allocation X . Since there is an equal number of agents and items, X will be a doubly stochastic matrix. According to the Birkhoff - von Neumann Theorem, every doubly stochastic matrix can be decomposed into a convex combination of some permutation matrices. Therefore, any probabilistic allocation X can be interpreted as a convex combination of a set of deterministic allocations. We denote the set of all possible instances by \mathcal{A} and denote the set of all possible allocation by \mathcal{X} . Given a mechanism M and a valuation profile $A \in \mathcal{A}$, as well as its allocation $X(A) \in \mathcal{X}$, we denote the expected utility of agent i by $u_i(X(A)) = \sum_j a_{ij}x_{ij}$ and denote the expected social welfare by $SW_M(X(A)) = \sum_i u_i(X(A))$. When the context is clear, we drop the allocation notation and simply refer them by $u_i(A)$ and $SW_M(A)$.

In Mechanism Design, agents are self-interested and may misreport their values if that results in a better matching (from their perspective). We are interested in *truthful mechanisms*, under which agents cannot improve their utilities by misreporting. Formally, $u_i(\mathbf{a}_i, \mathbf{a}_{-i}) \geq u_i(\mathbf{a}'_i, \mathbf{a}_{-i}), \forall i$, where \mathbf{a}_i is agent i 's true valuations, \mathbf{a}_{-i} is other agents' valuations, and \mathbf{a}'_i is any possible misreported valuation from agent i . The mechanism which is the focus of this paper, *Random Pri-*

ority (RP), is a truthful mechanism. It fixes an ordering of the agents uniformly at random and then allocates them their most preferred item from the set of available items based on this ordering.

The canonical measure of efficiency loss due to the restriction to the class of truthful mechanisms, compared to the optimal social welfare, is the *worst-case approximation ratio*,

$$r_{\text{worst}}(M) = \sup_{A \in \mathcal{A}} \frac{SW_{\text{OPT}}(A)}{SW_M(A)},$$

where $SW_{\text{OPT}}(A) = \max_{X \in \mathcal{X}} \sum_{i=1}^n u_i(X)$ is the optimal social welfare which is equal to the value of the maximum weight matching between agents and items. It is shown in [Filos-Ratsikas *et al.*, 2014] that Random Priority achieves the matching approximation ratio bound of $\Theta(\sqrt{n})$. The *average-case approximation ratio* of a truthful mechanism M is the *expectation of the ratio* of the optimal social welfare to the social welfare attained by mechanism M . That is,

$$r_{\text{average}}(M) = \mathbb{E}_{a_{ij} \sim D} \left[\frac{SW_{\text{OPT}}(A)}{SW_M(A)} \right],$$

where the valuation variable a_{ij} follows a distribution D . This notion of average-case analysis is a pointwise division that is in the same manner as the worst-case ratio $r_{\text{worst}}(M)$ and the smoothed ratio studied in [Deng *et al.*, 2017]. When randomly draw an instance from a distribution, it informs us the expected value of how far is the social welfare attainable by a truthful mechanism M on the instance compared to the optimal social welfare *on the same instance*.

In Bayesian mechanism design, the dominant approach is the *ratio of expectations*, defined as $\mathbb{E}[SW_{\text{OPT}}(A)] \leq r \cdot \mathbb{E}[SW_M(A)]$. When one's interest is the expected social welfare of a mechanism over all possible inputs compared to the expected optimal social welfare, rather than a pointwise comparison over the same instance, the ratio of expectations would be more appropriate. In most cases, it is more tractable, as it calculates two expectations separately.

3 Independent and Identically Distributed Random Values

Subject to the unit-range normalization, there are two preset values (0 and 1, respectively) in each row of a preference matrix A . Let S be the set of indices such that their corresponding entries in A are preset. That is, $S = \{(i, j) : a_{ij} = 0 \text{ or } a_{ij} = 1\}$. Obviously, $|S| = 2n$, where $|*|$ denotes the cardinality of a set¹. In this section, we allow the remaining $n^2 - 2n$ values to be independent and identically distributed random variables following a distribution D with expectation μ and variance σ^2 , i.e., $a_{ij} \sim D[0, 1]$, $\mathbb{E}_D[a_{ij}] = \mu$, $\text{Var}[a_{ij}] = \sigma^2, \forall i, j$. We will show that the average-case approximation ratio of RP is upper bounded by a constant, for any distribution $D[0, 1]$. In order to prove our main results, we partition the average-case ratio into two cases according to the social welfare attainable by RP. We

¹It does not matter which $2n$ of these valuations a_{ij} are preset. Our proofs hold for any choice of these $2n$ entries in A .

carefully choose a scalar such that in one case, the probability that the social welfare attainable by RP is smaller than the scalar is asymptotically small (even if there is a large efficiency loss in this case), and in the other case, the social welfare attainable by RP is asymptotically close to the optimal social welfare. En route to complying with these two conditions simultaneously, we employ the Central Limit Theorem and control the rate that the normalized sum of individual values converges to a normal distribution by carefully calibrating a scalar.

One of the tools that we use to control the convergence rate is the Berry-Esseen Theorem [Durrett, 2011].

Theorem 1 (Berry-Esseen) For $n \geq 1$, let X_1, \dots, X_n be i.i.d. random variables such that $\mathbb{E}[X_1] = 0$ and $\text{D}[X_1] = \sigma^2$. Denote $S_n = \frac{X_1 + \dots + X_n}{\sigma\sqrt{n}}$ and $F_n(x) = \Pr\{S_n \leq x\}$. Then

$$\sup_x |F_n(x) - \Phi(x)| \leq \frac{C \cdot \mathbb{E}[|X_1|^3]}{\sigma^3\sqrt{n}},$$

where $\Phi(x)$ is the cumulative distribution function of the standard Normal distribution, and $0.409 < C < 0.475$.

Now we will use the Berry-Esseen Theorem and the Central Limit Theorem to prove a key lemma.

Lemma 1 For a given preference instance A , the probability that the social welfare $SW_{RP}(A)$ attainable by Random Priority is less than λ , is bounded by the following inequality.

$$\Pr\{SW_{RP}(A) \leq \lambda\} \leq \frac{1}{2n\sqrt{\pi \ln n}} + \frac{C}{\sigma^3\sqrt{n(n-2)}},$$

where C is the constant in the Berry-Esseen Theorem and $\lambda = 1 + \mu(n-2) - \sigma\sqrt{\frac{2(n-2)}{n} \ln n}$.

Proof: Firstly, for any preference instance A , the social welfare attainable by Random Priority is lower bounded by the following inequality.

$$\begin{aligned} SW_{RP}(A) &\geq \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \\ &= \frac{1}{n} \sum_{(i,j) \in S} a_{ij} + \frac{1}{n} \sum_{(i,j) \notin S} a_{ij} \\ &= 1 + \frac{1}{n} \sum_{(i,j) \notin S} a_{ij}. \end{aligned}$$

Therefore,

$$\begin{aligned} \Pr\{SW_{RP}(A) \leq \lambda\} &\leq \Pr\left\{1 + \frac{1}{n} \sum_{(i,j) \notin S} a_{ij} \leq \lambda\right\} \\ &= \Pr\left\{\sum_{(i,j) \notin S} a_{ij} \leq \mu \cdot n(n-2) - \sigma\sqrt{2n(n-2) \ln n}\right\} \end{aligned} \quad (1)$$

Secondly, since a_{ij} are i.i.d. random variables and $\mathbb{E}[a_{ij}] = \mu$, $\text{Var}[a_{ij}] = \sigma^2, \forall i, j$, we know that $a_{ij} - \mu$ are

i.i.d. random variables such that $E[a_{ij} - \mu] = 0$, $\text{Var}[a_{ij} - \mu] = \sigma^2, \forall i, j$. Let

$$T_{n(n-2)} = \frac{\sum_{(i,j) \notin S} (a_{ij} - \mu)}{\sigma \sqrt{n(n-2)}}$$

and $G_{n(n-2)}(x) = \Pr\{T_{n(n-2)} \leq x\}$.

Following Theorem 1, the random variables $T_{n(n-2)}$ converge in distribution to the standard Normal distribution $N(0, 1)$. That is, $T_{n(n-2)} \xrightarrow{d} N(0, 1)$. In addition, the rate of convergence is bounded by

$$\begin{aligned} \sup_x |G_{n(n-2)}(x) - \Phi(x)| &\leq \frac{C \cdot E[|a_{ij} - \mu|]^3}{\sigma^3 \sqrt{n(n-2)}} \\ &\leq \frac{C}{\sigma^3 \sqrt{n(n-2)}}, \end{aligned}$$

where $(i, j) \notin S$, and the second inequality holds because $0 < a_{ij} < 1$ and $0 < \mu < 1$. Therefore,

$$\begin{aligned} |\Pr\{T_{n(n-2)} \leq x\} - \Phi(x)| &\leq \frac{C}{\sigma^3 \sqrt{n(n-2)}}, \quad \forall x. \\ \Rightarrow \left| \Pr\left\{ \frac{\sum_{(i,j) \notin S} (a_{ij} - \mu)}{\sigma \sqrt{n(n-2)}} \leq x \right\} - \Phi(x) \right| &\leq \frac{C}{\sigma^3 \sqrt{n(n-2)}} \\ \Rightarrow \Pr\left\{ \sum_{(i,j) \notin S} a_{ij} \leq \mu \cdot n(n-2) + x \cdot \sigma \sqrt{n(n-2)} \right\} \\ &\leq \int_{-\infty}^x \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt + \frac{C}{\sigma^3 \sqrt{n(n-2)}}. \end{aligned} \quad (2)$$

Let $x = -\sqrt{2 \ln n}$. Obviously, $-x \rightarrow +\infty$ when n approaches infinity. Following (2), we get that

$$\begin{aligned} \Pr\left\{ \sum_{(i,j) \notin S} a_{ij} \leq \mu \cdot n(n-2) - \sigma \sqrt{2 \ln n} \cdot \sqrt{n(n-2)} \right\} \\ &\leq \int_{-\infty}^{-\sqrt{2 \ln n}} \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt + \frac{C}{\sigma^3 \sqrt{n(n-2)}} \\ &= \int_{\sqrt{2 \ln n}}^{+\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt + \frac{C}{\sigma^3 \sqrt{n(n-2)}} \\ &\leq \frac{1}{\sqrt{2\pi} \sqrt{2 \ln n}} e^{-\ln n} + \frac{C}{\sigma^3 \sqrt{n(n-2)}} \\ &= \frac{1}{2n\sqrt{\pi \ln n}} + \frac{C}{\sigma^3 \sqrt{n(n-2)}}. \end{aligned} \quad (3)$$

Combining (1) and (3), we prove that

$$\Pr\{SW_{RP}(A) \leq \lambda\} \leq \frac{1}{2n\sqrt{\pi \ln n}} + \frac{C}{\sigma^3 \sqrt{n(n-2)}}.$$

□

We shall see that λ was chosen such that $\lambda \sim \Theta(n)$, and $\Pr\{SW_{RP}(A) \leq \lambda\} \sim O(\frac{1}{n})$. After establishing these building blocks, we are ready to prove the main theorem of this section.

Theorem 2 *The average-case approximation ratio of Random Priority, when agents' preferences are independent and identically distributed random variables, is upper bounded by the constant $1/\mu$, where μ is the expectation of these random variables.*

Proof: To calculate the expectation of the ratio of the optimal social welfare to the social welfare attained by RP, we partition the value space of the ratio into two cases according to the threshold parameter λ . In each case, we multiply the ratio by the probability that the case occurs. That is,

$$\begin{aligned} r_{\text{average}}(\text{RP}) &= E_{a_{ij} \sim D} \left[\frac{SW_{\text{OPT}}(A)}{SW_{\text{RP}}(A)} \right] \\ &\leq \Pr\{SW_{\text{RP}}(A) > \lambda\} \cdot \frac{n}{\lambda} \\ &\quad + \Pr\{SW_{\text{RP}}(A) \leq \lambda\} \cdot r_{\text{worst}} \end{aligned}$$

Here we plugged in the fact that $SW_{\text{OPT}}(A) \leq n$, for any instance A .

In addition, as shown in [Filos-Ratsikas *et al.*, 2014], $r_{\text{worst}} \sim \Theta(\sqrt{n})$, there exists a constant c_1 , such that $r_{\text{worst}} \leq c_1 \cdot \sqrt{n}$, for sufficiently large n . Also, it is obvious that $\Pr\{SW_{\text{RP}}(A) > \lambda\} \leq 1$. Now, we plug in these fact and the results established in Lemma 1.

$$\begin{aligned} r_{\text{average}}(\text{RP}) &\leq 1 \cdot \frac{n}{1 + \mu(n-2) - \sigma \sqrt{\frac{2(n-2)}{n} \ln n}} \\ &\quad + \left[\frac{1}{2n\sqrt{\pi \ln n}} + \frac{C}{\sigma^3 \sqrt{n(n-2)}} \right] \cdot c_1 \sqrt{n} \\ &\leq \frac{1}{\mu} \left(\frac{1}{1 - \frac{2}{n} + \frac{1}{\mu n} - \frac{\sigma \sqrt{2 \ln n}}{\mu n}} \right) \\ &\quad + \frac{c_1}{2\sqrt{\pi n \ln n}} + \frac{C \cdot c_1}{\sigma^3 \sqrt{n-2}} \end{aligned}$$

Since $\frac{1}{1-|t|} < 1 + 2|t|$, when $|t| < \frac{1}{2}$. As long as $\left| \frac{2}{n} - \frac{1}{\mu n} + \frac{\sigma \sqrt{2 \ln n}}{\mu n} \right| < \frac{1}{2}$, we conclude that

$$\begin{aligned} r_{\text{average}}(\text{RP}) &\leq \frac{1}{\mu} \left(1 + 2 \left| \frac{2}{n} - \frac{1}{\mu n} + \frac{\sigma \sqrt{2 \ln n}}{\mu n} \right| \right) \\ &\quad + \frac{c_1}{2\sqrt{\pi n \ln n}} + \frac{C \cdot c_1}{\sigma^3 \sqrt{n-2}} \\ &= \frac{1}{\mu} + \frac{2}{\mu} \left| \frac{2}{n} - \frac{1}{\mu n} + \frac{\sigma \sqrt{2 \ln n}}{\mu n} \right| \\ &\quad + \frac{c_1}{2\sqrt{\pi n \ln n}} + \frac{C \cdot c_1}{\sigma^3 \sqrt{n-2}} \\ &\rightarrow \frac{1}{\mu}. \end{aligned}$$

□

4 Independent but not Necessarily Identical Random Values

In this section, we further generalize the results in Section 3 to the scenario that the $n^2 - 2n$ elements in $\{a_{ij} | (i, j) \notin S\}$ are independent random values but are not restricted to following an identical distribution. So, they may have different expectations $E[a_{ij}] = \mu_{ij}$ and variance $\text{Var}[a_{ij}] = \sigma_{ij}^2, \forall i, j$. We will show that under mild conditions, the average-case approximation ratio of RP has a constant upper bound. The primary obstruction in this generalization is the difficulty to appropriately control the rate of convergence of the normalized sum of the random values, in order to obtain a constant upper bound. However, we are able to pinpoint a mild condition, such that when the random variables a_{ij} comply with the condition, we can establish similar building blocks to the last section. By carefully calibrating the parameter λ , we can assemble them in a compatible way to get a constant upper bound.

Firstly, we identify the following mild conditions:

(i) $\sum_{(i,j) \notin S} \mu_{ij} = \Omega(n^2)$; (ii) $\sum_{(i,j) \notin S} \sigma_{ij}^2 = \omega(n)$. Note that there are $n^2 - 2n$ elements in the set $\{a_{ij} | (i, j) \notin S\}$, so the first condition simply implies that there would not be many of these a_{ij} whose expectations are asymptotically small. The second condition merely implies that the variances of the valuations are not too small. In other words, the first condition excludes those instances that many agents' preferences are negligible; the second condition requires their preferences to admit a magnitude of variation.

Secondly, to control the rate of convergence, we will employ the following theorem which is a refined version of the Berry-Esseen Theorem. It holds for non-identically distributed random variables [Esseen, 1945].

Theorem 3 Let Z_1, \dots, Z_n be independent random variables such that $E[Z_i] = 0$, $D[Z_i] = \sigma_i^2$, and $E[|Z_i|^3] < \infty$. Denote $X_n = \frac{Z_1 + \dots + Z_n}{\sqrt{\sum_i \sigma_i^2}}$ and $F_n(x) = \Pr\{X_n \leq x\}$. Then there exists a constant C' such that

$$\sup_x |F_n(x) - \Phi(x)| \leq \frac{C' \cdot \sum_i E[|Z_i|^3]}{(\sum_i \sigma_i^2)^{\frac{3}{2}}},$$

where $\Phi(x)$ is the CDF of the standard Normal distribution.

Now, let $Z_{ij} := a_{ij} - \mu_{ij}$, where $(i, j) \notin S$. Denote their normalized sum of Z_{ij} and the CDF of $X_{n(n-2)}$ by

$$X_{n(n-2)} := \frac{\sum_{i,j} Z_{ij}}{\sqrt{\sum_{i,j} \sigma_{ij}^2}}, \quad F_{X_{n(n-2)}}(x) = \Pr(X_{n(n-2)} \leq x).$$

Next, let $\lambda = 1 + \frac{\sum_{(i,j) \notin S} \mu_{ij}}{n} - \frac{\sqrt{2 \ln n}}{n} \cdot \sqrt{\sum_{(i,j) \notin S} \sigma_{ij}^2}$.

With this carefully chosen value of λ , we are able to show the following lemma.

Lemma 2 For a given preference matrix A , the probability that the social welfare $SW_{RP}(A)$ attainable by Random Priority is less than λ , is bounded by the following inequality.

$$\Pr\{SW_{RP}(A) \leq \lambda\} \leq \frac{1}{2n\sqrt{\pi \ln n}} + \frac{C'}{\sqrt{\sum_{i,j} \sigma_{ij}^2}},$$

where C' is the constant in Theorem 3.

Proof: On the one hand, $SW_{RP}(A) \geq 1 + \frac{1}{n} \sum_{(i,j) \notin S} a_{ij}$ implies that

$$\begin{aligned} \Pr\{SW_{RP}(A) \leq \lambda\} &\leq \Pr\left\{1 + \frac{1}{n} \sum_{(i,j) \notin S} a_{ij} \leq \lambda\right\} \\ &= \Pr\left\{\sum_{(i,j) \notin S} a_{ij} \leq \sum_{(i,j) \notin S} \mu_{ij} - \sqrt{2 \ln n} \cdot \sqrt{\sum_{(i,j) \notin S} \sigma_{ij}^2}\right\} \end{aligned} \quad (4)$$

On the other hand, according to Theorem 3, there exists a constant C' , such that

$$\sup_x |F_{X_{n(n-2)}}(x) - \Phi(x)| \leq \frac{C' \cdot \sum_{i,j} E[|Z_{ij}|^3]}{(\sum_{i,j} \sigma_{ij}^2)^{\frac{3}{2}}}.$$

Since $|Z_{ij}| \leq 1$, we have $E[|Z_{ij}|^3] \leq E[|Z_{ij}|^2] = \sigma_{ij}^2$. So,

$$\sup_x |F_{X_{n(n-2)}}(x) - \Phi(x)| \leq \frac{C'}{\sqrt{\sum_{i,j} \sigma_{ij}^2}}.$$

Therefore, $\forall x$,

$$\begin{aligned} |\Pr\{X_{n(n-2)} \leq x\} - \Phi(x)| &\leq \frac{C'}{\sqrt{\sum_{i,j} \sigma_{ij}^2}}, \\ \Rightarrow \left| \Pr\left\{\frac{\sum_{i,j} (a_{ij} - \mu_{ij})}{\sqrt{\sum_{i,j} \sigma_{ij}^2}} \leq x\right\} - \Phi(x) \right| &\leq \frac{C'}{\sqrt{\sum_{i,j} \sigma_{ij}^2}} \\ \Rightarrow \Pr\left\{\sum_{i,j} a_{ij} \leq \sum_{i,j} \mu_{ij} + x \cdot \sqrt{\sum_{i,j} \sigma_{ij}^2}\right\} & \\ \leq \int_{-\infty}^x \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt + \frac{C'}{\sqrt{\sum_{i,j} \sigma_{ij}^2}}. & \end{aligned} \quad (5)$$

Let $x = -\sqrt{2 \ln n}$. Obviously, $-x \rightarrow +\infty$ when n approaches infinity. Following (5), we obtain that

$$\begin{aligned} \Pr\left\{\sum_{i,j} a_{ij} \leq \sum_{i,j} \mu_{ij} - \sqrt{2 \ln n} \cdot \sqrt{\sum_{i,j} \sigma_{ij}^2}\right\} & \\ \leq \int_{-\infty}^{-\sqrt{2 \ln n}} \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} dt + \frac{C'}{\sqrt{\sum_{i,j} \sigma_{ij}^2}} & \\ \leq \frac{1}{\sqrt{2\pi} \sqrt{2 \ln n}} e^{-\ln n} + \frac{C'}{\sqrt{\sum_{i,j} \sigma_{ij}^2}} & \\ \leq \frac{1}{2n\sqrt{\pi \ln n}} + \frac{C'}{\sqrt{\sum_{i,j} \sigma_{ij}^2}}. & \end{aligned} \quad (6)$$

Combining (4) and (6), we complete the proof. \square

Finally, we are ready to prove the main theorem of this section.

Theorem 4 The average-case approximation ratio of Random Priority, when agents' preferences are independent but not necessarily identically distributed random variables, is upper bounded by a constant.

Proof: We partition the value space of the ratio into two cases according to the threshold parameter λ . In each case, we multiply the ratio by the probability that the case occurs. That is,

$$\begin{aligned} r_{\text{average}}(\text{RP}) &= \mathbb{E}_{a_{ij} \sim D_{ij}} \left[\frac{SW_{\text{OPT}}(A)}{SW_{\text{RP}}(A)} \right] \\ &\leq \Pr\{SW_{\text{RP}}(A) > \lambda\} \cdot \frac{n}{\lambda} + \Pr\{SW_{\text{RP}}(A) \leq \lambda\} \cdot c_1 \sqrt{n} \end{aligned}$$

We assemble the above building blocks and get that

$$\begin{aligned} r_{\text{average}}(\text{RP}) &\leq 1 \cdot \frac{n}{1 + \frac{\sum_{i,j} \mu_{ij}}{n} - \frac{\sqrt{2 \ln n}}{n} \cdot \sqrt{\sum_{i,j} \sigma_{ij}^2}} \\ &\quad + \left(\frac{1}{2n\sqrt{\pi \ln n}} + \frac{C'}{\sqrt{\sum_{i,j} \sigma_{ij}^2}} \right) \cdot c_1 \sqrt{n} \\ &\leq \frac{n^2}{\sum \mu_{i,j}} \left(\frac{1}{1 - \left(\sqrt{2 \ln n} \cdot \frac{\sqrt{\sum \sigma_{ij}^2}}{\sum \mu_{ij}} - \frac{n}{\sum \mu_{ij}} \right)} \right) \\ &\quad + \frac{c_1}{2\sqrt{\pi n \ln n}} + \frac{C' \cdot c_1 \sqrt{n}}{\sqrt{\sum \sigma_{ij}^2}} \end{aligned}$$

According to the first condition, there $\exists \underline{\mu}$, $\exists N_0$, such that $\forall n > N_0$, $\frac{n^2}{\sum \mu_{ij}} \leq \frac{1}{\underline{\mu}}$. Note that $0 < a_{ij} < 1$ and $0 < \mu_{ij} < 1$, so $\sum \mu_{ij} < n^2$, hence $\frac{1}{\underline{\mu}} > 1$. We also know that $\sum \sigma_{ij}^2 < n^2$. Together with the first condition, we have that $\sqrt{2 \ln n} \cdot \frac{\sqrt{\sum \sigma_{ij}^2}}{\sum \mu_{ij}} \rightarrow 0$ and $\frac{n}{\sum \mu_{ij}} \rightarrow 0$, when $n \rightarrow \infty$. Therefore, there $\exists N_1$, such that when $n > N_1$, $\left| \frac{\sqrt{2 \ln n} \sqrt{\sum \sigma_{ij}^2}}{\sum \mu_{ij}} - \frac{n}{\sum \mu_{ij}} \right| < \frac{1}{2}$. Hence, when $n > \max\{N_0, N_1\}$, we have that

$$\begin{aligned} r_{\text{average}}(\text{RP}) &\leq \frac{1}{\underline{\mu}} \left(1 + 2 \left| \frac{\sqrt{2 \ln n} \sqrt{\sum \sigma_{ij}^2}}{\sum \mu_{ij}} - \frac{n}{\sum \mu_{ij}} \right| \right) \\ &\quad + \frac{c_1}{2\sqrt{\pi n \ln n}} + \frac{C' \cdot c_1 \sqrt{n}}{\sqrt{\sum \sigma_{ij}^2}} \\ &= \frac{1}{\underline{\mu}} + \frac{2}{\underline{\mu}} \left| \frac{\sqrt{2 \ln n} \sqrt{\sum \sigma_{ij}^2}}{\sum \mu_{ij}} - \frac{n}{\sum \mu_{ij}} \right| \\ &\quad + \frac{c_1}{2\sqrt{\pi n \ln n}} + \frac{C' \cdot c_1 \sqrt{n}}{\sqrt{\sum \sigma_{ij}^2}} \end{aligned}$$

Because of condition (ii), we know that $\frac{\sqrt{n}}{\sqrt{\sum \sigma_{ij}^2}} \rightarrow 0$. In conclusion,

$$r_{\text{average}}(\text{RP}) \rightarrow \frac{1}{\underline{\mu}} + 0 = \frac{1}{\underline{\mu}}.$$

□

5 Conclusion

This paper extended the average-case analysis in [Deng *et al.*, 2017] from a uniform distribution to any independent distribution and showed a constant upper bound of the approximation ratio. The average-case analysis complements classical worst-case analysis when the worst-case performance is insufficient to characterize the performance of a mechanism. Our results further justify the wide-applicability of the Random Priority mechanism.

There are a few technical points we would like to highlight here. Firstly, the techniques presented in this paper are probably applicable to analyzing other mechanisms and other domains. Secondly, there are various generalizations of the Berry-Esseen Theorem, and each of them may be cast to prove similar results in Section 4. In an earlier version of the present paper, we had also independently proved a version of the convergence rate. The difference in utilizing different versions of these theorems is that they each require a set of conditions to make the rest of the proof work, and the interpretations of those conditions could be different.

There are a number of problems remain open as well. For example, one may be interested in investigating the average-case ratio of Random Priority in correlated domains. Also, it would be interesting to proving tighter bounds by making more use of the structure of the assignment problem domain.

References

- [Abdulkadiroğlu and Sönmez, 2013] Atila Abdulkadiroğlu and Tayfun Sönmez. Matching Markets: Theory and Practice. *Advances in Economics and Econometrics (Tenth World Congress)*, pages 3–47, 2013.
- [Abdulkadiroglu and Sönmez, 1998] Atila Abdulkadiroglu and Tayfun Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–702, 1998.
- [Abdulkadiroglu and Sönmez, 1999] Atila Abdulkadiroglu and Tayfun Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, 1999.
- [Anshelevich *et al.*, 2018] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. Approximating optimal social choice under metric preferences. *Artif. Intell.*, 264:27–51, 2018.
- [Anshelevich, 2016] Elliot Anshelevich. Ordinal approximation in matching and social choice. *SIGecom Exchanges*, 15(1):60–64, 2016.
- [Archer *et al.*, 2003] Aaron Archer, Christos H. Papadimitriou, Kunal Talwar, and Éva Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. In *SODA*, pages 205–214, 2003.
- [Aziz *et al.*, 2013] Haris Aziz, Felix Brandt, and Markus Brill. The computational complexity of random serial dictatorship. In *WINE*, pages 24–25, 2013.
- [Barbera, 2010] Salvador Barbera. Strategy-proof Social Choice. In K. J. Arrow, A. K. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 2, chapter 25. North-Holland: Amsterdam, 2010.

- [Bhalgat *et al.*, 2011] Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. Social welfare in one-sided matching markets without money. In *APPROX-RANDOM*, pages 87–98, 2011.
- [Bogomolnaia and Moulin, 2001] Anna Bogomolnaia and Hervé Moulin. A New Solution to the Random Assignment Problem. *Journal of Economic Theory*, 100:295–328, 2001.
- [Chawla and Sivan, 2014] Shuchi Chawla and Balasubramanian Sivan. Bayesian algorithmic mechanism design. *SIGecom Exchanges*, 13(1):5–49, 2014.
- [Chen *et al.*, 2013] Yiling Chen, John Lai, David C. Parkes, and Ariel D. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1):284–297, 2013.
- [Christodoulou *et al.*, 2009] G. Christodoulou, E. Koutsoupias, and A. Vidali. A lower bound for scheduling mechanisms. *Algorithmica*, 55(4):729–740, 2009.
- [Christodoulou *et al.*, 2016] George Christodoulou, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Paul W. Goldberg, Jie Zhang, and Jinshan Zhang. Social welfare in one-sided matching mechanisms. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1297–1298, 2016.
- [Deng *et al.*, 2017] Xiaotie Deng, Yansong Gao, and Jie Zhang. Smoothed and average-case approximation ratios of mechanisms: Beyond the worst-case analysis. In *42nd International Symposium on Mathematical Foundations of Computer Science, MFCS*, pages 16:1–16:15, 2017.
- [Dughmi and Ghosh, 2010] Shaddin Dughmi and Arpita Ghosh. Truthful assignment without money. In *ACM Conference on Electronic Commerce*, pages 325–334, 2010.
- [Durrett, 2011] Rick Durrett. *Probability: Theory and examples*. Cambridge University Press, 2011.
- [Esseen, 1945] Carl-Gustav Esseen. Fourier analysis of distribution functions. a mathematical study of the laplace-gaussian law. *Acta Mathematica*, 77:1–125, 1945.
- [Filos-Ratsikas *et al.*, 2014] Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *Algorithmic Game Theory, SAGT 2014*, pages 1–12, 2014.
- [Filos-Ratsikas *et al.*, 2015] Aris Filos-Ratsikas, Minming Li, Jie Zhang, and Qiang Zhang. Facility location with double-peaked preferences. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 893–899, 2015.
- [Hajaj *et al.*, 2015] Chen Hajaj, John P. Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne. Strategy-proof and efficient kidney exchange using a credit mechanism. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 921–928, 2015.
- [Hartline and Lucier, 2010] Jason D. Hartline and Brendan Lucier. Bayesian algorithmic mechanism design. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010*, pages 301–310, 2010.
- [Hartline and Roughgarden, 2009] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *ACM-EC*, pages 225–234, 2009.
- [Hylland and Zeckhauser, 1979] Aanund Hylland and Richard Zeckhauser. The Efficient Allocation of Individuals to Positions. *The Journal of Political Economy*, 87(2):293–314, 1979.
- [Katta and Sethuraman, 2006] Akshay-Kumar Katta and Jay Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic theory*, 131(1):231–250, 2006.
- [Kesten, 2009] Onur Kesten. Why do popular mechanisms lack efficiency in random environments? *Journal of Economic Theory*, 144:2209–2226, 2009.
- [Koutsoupias, 2014] Elias Koutsoupias. Scheduling without payments. *Theory Comput. Syst.*, 54(3):375–387, 2014.
- [Mennle and Seuken, 2014] Timo Mennle and Sven Seuken. An axiomatic approach to characterizing and relaxing strategyproofness of one-sided matching mechanisms. In *ACM-EC*, pages 37–38, 2014.
- [Mu’alem and Nisan, 2008] Ahuva Mu’alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64(2):612–631, 2008.
- [Nisan and Ronen, 1999] Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing (STOC)*, pages 129–140, 1999.
- [Nisan *et al.*, 2007] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [Pathak and Sethuraman, 2011] P. A. Pathak and J. Sethuraman. Lotteries in student assignment: An equivalence result. *Theoretical Economics*, 6, 2011.
- [Procaccia and Tennenholtz, 2009] Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *ACM-EC*, pages 177–186, 2009.
- [Roth and Sotomayor, 1990] A. E. Roth and M. Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, 1990.
- [Sönmez and Ünver, 2011] Tayfun Sönmez and Utku Ünver. Matching, allocation and exchange of discrete resources. *Handbook of Social Economics*, 1A:781–852, 2011.
- [Svensson, 1999] Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.
- [Zhang, 2018] Jie Zhang. Average-case approximation ratio of scheduling without payments. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, pages 1298–1305, 2018.
- [Zhou, 1990] Lin Zhou. On a Conjecture by Gale about One-Sided Matching Problems. *Journal of Economic Theory*, 52:123–135, 1990.