

# Differentially Private Fair Division

Pasin Manurangsi<sup>1</sup>, Warut Suksompong<sup>2</sup>

<sup>1</sup>Google Research

<sup>2</sup>National University of Singapore

## Abstract

Fairness and privacy are two important concerns in social decision-making processes such as resource allocation. We study privacy in the fair allocation of indivisible resources using the well-established framework of differential privacy. We present algorithms for approximate envy-freeness and proportionality when two instances are considered to be adjacent if they differ only on the utility of a single agent for a single item. On the other hand, we provide strong negative results for both fairness criteria when the adjacency notion allows the entire utility function of a single agent to change.

## 1 Introduction

Fairness is a principal concern in numerous social decision-making processes, not least when it comes to allocating scarce resources among interested parties. Whether we divide equipment between healthcare personnel, assign facility time slots to potential users, or distribute office space among working groups in an organization, it is desirable that all parties involved feel fairly treated. While fair division has been studied in economics for several decades (Brams and Taylor 1996; Moulin 2003, 2019), the subject has received substantial interest from computer scientists in recent years, much of which has concentrated on the fair allocation of indivisible resources (Bouveret, Chevaleyre, and Maudet 2016; Markakis 2017; Walsh 2020; Suksompong 2021; Amanatidis et al. 2022; Aziz et al. 2022).

The fair division literature typically focuses on satisfying concrete fairness criteria. Two of the most prominent criteria are *envy-freeness* and *proportionality*. In an envy-free allocation, no agent prefers to have another agent’s bundle instead of her own. In a proportional allocation, every agent receives value at least  $1/n$  of her value for the entire resource, where  $n$  denotes the number of agents. As neither of these criteria is always satisfiable,<sup>1</sup> researchers have proposed the relaxations *envy-freeness up to  $c$  items* (*EF $c$* ) and *proportionality up to  $c$  items* (*PROP $c$* ); here,  $c$  is a non-negative integer parameter. Under additive utilities, *EF $c$*  implies *PROP $c$*  for every  $c$ ,<sup>2</sup> and an *EF1* allocation (which must

also be *PROP1*) is guaranteed to exist (Lipton et al. 2004).

In addition to fairness, another consideration that has become increasingly important nowadays—as large amounts of data are constantly collected, processed, and analyzed—is privacy. Indeed, an agent participating in a resource allocation procedure may not want other participants to know her preferences if she considers these as sensitive information, for example, if these preferences correspond to the times when she is available to use a facility, or if they represent her valuations for potential team members when distributing employees in an organization. Consequently, a desirable procedure should ensure that an individual participant’s preferences cannot be inferred based on the output of the procedure. Achieving privacy alone is trivial, as the procedure can simply ignore the agents’ preferences and always output a fixed allocation that it announces publicly in advance. However, it is clear that such a procedure can be highly unfair for certain preferences of the agents. Despite its significance, the issue of privacy has been largely unaddressed in the fair division literature as far as we are aware.<sup>3</sup>

In this paper, we investigate the fundamental question of whether fairness and privacy can be attained simultaneously in the allocation of indivisible resources. We use the well-established framework of *differential privacy* (*DP*), which has been widely adopted not only in academia but also in industry (Erlingsson, Pihur, and Korolova 2014; Shankland 2014; Greenberg 2016; Apple Differential Privacy Team 2017; Ding, Kulkarni, and Yekhanin 2017) as well as government sectors (Abowd 2018). Intuitively, the output distribution of a (randomized) *DP*<sup>4</sup> algorithm should not change by much when a single “entry” of the input is modified. *DP* provides a privacy protection for individual entries by ensuring that an adversary with access to the output can only gain limited information about each individual entry. At the same time, *DP* algorithms often still provide useful outputs based on the aggregated information. We outline the tools and concepts from *DP* used in this work in Section 2.2.<sup>5</sup>

<sup>3</sup>Sun and Yang (2009) studied a setting in which agents have “reservation values” over items, and considered a mechanism to be private if it does not require agents to reveal these values.

<sup>4</sup>We use the abbreviation *DP* for both “differential privacy” and “differentially private”.

<sup>5</sup>For in-depth treatments of the subject, we refer to the surveys by Dwork (2008) and Dwork and Roth (2014).

<sup>1</sup>For example, imagine two siblings fighting for one toy.

<sup>2</sup>See the proof of Proposition 2.2(b) in our prior work (Manurangsi and Suksompong 2022a).

		Agent-Level DP	(Agent $\times$ Item)-Level DP
EF	Upper	$O(m/n)$ (Trivial)	$O(n \log m)$ (Theorem 4.1)
	Lower	$\Omega(\sqrt{\frac{m}{n} \log n})$ (Theorem 3.1)	$\Omega(\log m)$ (Theorem 4.8)
PROP	Upper	$O(m/n)$ (Trivial)	$O(\log m)$ (Theorem 4.7)
	Lower	$\Omega(\sqrt{\frac{m}{n}})$ (Theorem 3.2)	$\Omega((\log m)/n)$ (Theorem 4.10)

Table 1: Overview of our results. We display the smallest  $c$  such that there is an  $\varepsilon$ -DP algorithm that outputs an EF $_c$  or PROP $_c$  allocation with probability at least  $1 - \beta$ . For simplicity of the bounds, we assume that  $m > n^2$ ,  $\varepsilon$  is a small constant, and  $\beta$  is sufficiently small (depending only on  $\varepsilon, n$  and not on  $m$ ). All upper bounds hold even for connected allocations. Lower bounds for (agent  $\times$  item)-level DP hold only for connected allocations, but those for agent-level DP hold for arbitrary allocations. The bounds  $O(m/n)$  follow trivially from outputting a fixed allocation in which each agent receives  $O(m/n)$  items.

As alluded to above, DP is defined with respect to what one considers to be an entry of the input, or equivalently, in terms of adjacent inputs. We consider two notions of adjacency between fair division instances. For *agent-level adjacency*, two instances are considered to be adjacent if they differ on the utility function of at most one agent. For *(agent  $\times$  item)-level adjacency*, two instances are adjacent if they differ at most on the utility of a single agent for a single item. We work with the standard definition of  $\varepsilon$ -differential privacy ( $\varepsilon$ -DP): for a parameter  $\varepsilon \geq 0$ , an algorithm is said to satisfy  $\varepsilon$ -DP if the probability that it outputs a certain allocation for an input and the corresponding probability for an adjacent input differ by a factor of at most  $e^\varepsilon$ . Note that, for the same  $\varepsilon$ , *agent-level DP* offers a stronger privacy protection for an entire utility function of an individual agent, whereas *(agent  $\times$  item)-level DP* only offers a protection for a utility of a single agent for a single item. Our goal is to devise  $\varepsilon$ -DP algorithms that output an EF $_c$  or PROP $_c$  allocation for a small value of  $c$  with sufficiently high probability, or to prove that this task is impossible. Denote by  $n$  and  $m$  the number of agents and items, respectively.

We begin in Section 3 by considering agent-level DP. For this demanding benchmark, we establish strong lower bounds with respect to both approximate envy-freeness and proportionality (Theorems 3.1 and 3.2). In both cases, our lower bounds imply that, for fixed  $n$  and  $\varepsilon$ , no  $\varepsilon$ -DP algorithm can output an EF $_c$  or PROP $_c$  allocation for  $c = o(\sqrt{m})$  with some large constant probability. Our results hold even when the agents have binary additive utilities, and indicate that agent-level DP is too stringent to permit algorithms with significant fairness guarantees.

In Section 4, we turn our attention to (agent  $\times$  item)-level DP, and deliver encouraging news for this more relaxed notion. In contrast to the previous lower bounds, we present  $\varepsilon$ -DP algorithms for EF $_c$  and PROP $_c$  where  $c$  only grows logarithmically in  $m$  for fixed  $n$  and  $\varepsilon$  (Theorems 4.1 and 4.7). Our EF $_c$  algorithm works for arbitrary monotone utility functions, whereas our PROP $_c$  algorithm allows (not necessarily binary) additive utilities. Moreover, our algorithms always output allocations that are *connected*.<sup>6</sup> We complement these results by showing a tight

<sup>6</sup>See Section 2 for the definition. Connectivity can be desirable when there is a spatial or temporal order of the items, for instance,

lower bound of  $\Omega(\log m)$  for connected allocations (Theorems 4.8 and 4.10), even with binary additive utilities.

A summary of our results can be found in Table 1.

## 2 Preliminaries

### 2.1 Fair Division

In fair division of indivisible items, there is a set  $N = [n]$  of agents and a set  $M = [m]$  of items, where  $[k]$  denotes the set  $\{1, 2, \dots, k\}$  for each positive integer  $k$ . The utility function of agent  $i$  is given by  $u_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ . Throughout this work, we assume that utility functions are *monotone*, that is,  $u_i(S) \leq u_i(T)$  for any  $S \subseteq T \subseteq M$ . For a single item  $j \in M$ , we write  $u_i(j)$  instead of  $u_i(\{j\})$ . We seek to output an *allocation*  $A = (A_1, \dots, A_n)$ , which is an ordered partition of  $M$  into  $n$  bundles.

We consider two important fairness notions. Let  $c$  be a non-negative integer.

- An allocation  $A$  is said to be *envy-free up to  $c$  items (EF $_c$ )* if, for any  $i, i' \in N$ , there exists  $S \subseteq A_{i'}$  with  $|S| \leq c$  such that  $u_i(A_i) \geq u_i(A_{i'} \setminus S)$ .
- An allocation  $A$  is said to be *proportional up to  $c$  items (PROP $_c$ )* if, for any  $i \in N$ , there exists  $S \subseteq M \setminus A_i$  with  $|S| \leq c$  such that  $u_i(A_i) \geq u_i(M/n) - u_i(S)$ .

We say that an allocation  $A$  is *connected* if each  $A_i$  corresponds to an interval, i.e.,  $A_i = \{\ell, \ell + 1, \dots, r\}$  for some  $\ell, r \in M$ ; it is possible that  $A_i$  is a singleton or empty. Let  $\mathcal{P}^{\text{conn}}(m, n)$  denote the set of all connected allocations. We will use the following result on the existence of connected EF2 allocations.<sup>7</sup>

**Theorem 2.1** ((Bilò et al. 2022)). *For any  $m, n \in \mathbb{N}$  and any monotone utility functions, there exists a connected EF2 allocation.*

A utility function  $u_i$  is *additive* if  $u_i(S) = \sum_{j \in S} u_i(j)$  for all  $S \subseteq M$ . Furthermore, an additive utility function is said to be *binary* if  $u_i(j) \in \{0, 1\}$  for all  $j \in M$ .

when allocating time slots to facility users or offices along a corridor to research groups (Bouveret et al. 2017; Suksompong 2019; Bei et al. 2022; Bilò et al. 2022).

<sup>7</sup>Recently, Igarashi (2023) improved this guarantee to EF1.

## 2.2 Differential Privacy

Let us start by recalling the general definition of DP. Denote by  $\mathcal{X}$  the set of all possible inputs to the algorithm.

**Definition 2.2** (Differential Privacy (Dwork et al. 2006b)). Let  $\varepsilon \geq 0$  be a non-negative real number. A randomized algorithm  $\mathcal{M}$  is said to be  $\varepsilon$ -differentially private ( $\varepsilon$ -DP) if, for every pair of adjacent inputs  $X, X'$ , it holds that

$$\Pr[\mathcal{M}(X) = o] \leq e^\varepsilon \cdot \Pr[\mathcal{M}(X') = o]$$

for all  $o \in \text{range}(\mathcal{M})$ .

An input in our fair division context consists of the agents' utility functions. Different notions of adjacency lead to different levels of privacy protection. We consider two natural notions: agent-level DP and (agent  $\times$  item)-level DP.

**Definition 2.3** (Agent-Level DP). Two inputs  $(u_i)_{i \in N}$  and  $(u'_i)_{i \in N}$  are said to be *agent-level adjacent* if they coincide on all but a single agent, i.e., there exists  $i^* \in N$  such that  $u_i = u'_i$  for all  $i \in N \setminus \{i^*\}$ .

An algorithm that is  $\varepsilon$ -DP against this adjacency notion is said to be *agent-level  $\varepsilon$ -DP*.

**Definition 2.4** ((Agent  $\times$  Item)-Level DP). Two inputs  $(u_i)_{i \in N}$  and  $(u'_i)_{i \in N}$  are said to be *(agent  $\times$  item)-level adjacent* if they coincide on all but the utility of a single agent for a single item, i.e., there exist  $i^* \in N, j^* \in M$  such that

- $u_i = u'_i$  for all  $i \in N \setminus \{i^*\}$ , and
- $u_{i^*}(S) = u'_{i^*}(S)$  for all  $S \subseteq M \setminus \{j^*\}$ .

An algorithm that is  $\varepsilon$ -DP against this adjacency notion is said to be *(agent  $\times$  item)-level  $\varepsilon$ -DP*.

It is clear that agent-level DP is a more demanding notion than (agent  $\times$  item)-level DP. Specifically, the former provides a stronger privacy protection than the latter, and designing an algorithm for the former is more difficult. Indeed, we will prove strong lower bounds for agent-level DP and present algorithms for (agent  $\times$  item)-level DP.

Next, we outline several tools from the DP literature that will be useful for our proofs.

**Basic Composition.** The first tool that we will use is the composition of DP: the result of running multiple DP algorithms remains DP, but with a worse privacy parameter.

**Theorem 2.5** (Basic Composition of DP, e.g., (Dwork and Roth 2014)). *An algorithm that is a result of running two algorithms (possibly in an adaptive manner) that are  $\varepsilon_1$ -DP and  $\varepsilon_2$ -DP, respectively, is  $(\varepsilon_1 + \varepsilon_2)$ -DP.*

**Group Privacy.** While differential privacy offers protection primarily against the adjacency notion for which it is defined, it also offers protection against more general adjacency notions. Below we state one such protection, which is often referred to as *group differential privacy*.

Let  $\sim$  be any adjacency relation. For  $k \in \mathbb{N}$ , let us define  $\sim_k$  as the adjacency relation where  $X \sim_k X'$  if and only if there exists a sequence  $X_0, \dots, X_k$  such that  $X_0 = X, X_k = X'$ , and  $X_{i-1} \sim X_i$  for all  $i \in [k]$ .

**Lemma 2.6** (Group Differential Privacy, e.g., (Vadhan 2017)). *Let  $k \in \mathbb{N}$ . If an algorithm is  $\varepsilon$ -DP with respect to an adjacency notion  $\sim$ , it is also  $(k\varepsilon)$ -DP with respect to the adjacency notion  $\sim_k$ .*

As an immediate consequence of Lemma 2.6, any (agent  $\times$  item)-level  $\varepsilon$ -DP algorithm is also agent-level  $(m\varepsilon)$ -DP. However, the factor  $m\varepsilon$  makes the latter guarantee rather weak, especially as  $m$  grows.

**Sensitivity.** We next define the *sensitivity* of a function, which will be used multiple times in this work. Note that the definition depends on the adjacency notion, but we do not explicitly include it in the notation for convenience.

**Definition 2.7.** The *sensitivity* of a function  $f : \mathcal{X} \rightarrow \mathbb{R}$  (with respect to adjacency notion  $\sim$ ) is defined as  $\Delta(f) := \max_{X \sim X'} |f(X) - f(X')|$ .

Sensitivity is a key notion in DP. As shown by Dwork et al. (2006b), the *Laplace mechanism*—which outputs  $f(X) + Z$  where  $Z$  is drawn from the Laplace distribution<sup>8</sup> with scale  $(\Delta(f)/\varepsilon)$ —satisfies  $\varepsilon$ -DP. This means that if a function has low sensitivity, then we can estimate it to within a small error (with high probability).

**Sparse Vector Technique.** We will use the so-called *sparse vector technique (SVT)*. The setting is that there are low-sensitivity functions  $f_1, \dots, f_H : \mathcal{X} \rightarrow \mathbb{R}$ . We want to find the first function  $f_i$  whose value is above a target threshold. A straightforward approach would be to add Laplace noise to each  $f_i(X)$  and then select the function accordingly; due to the basic composition theorem, this would require us to add noise of scale  $O(H/\varepsilon)$  to each function. SVT allows us to reduce the dependency on  $H$  to merely  $O(\log H)$ . The technique was first introduced by Dwork et al. (2009), and the convenient version below is due to Dwork and Roth (2014, Theorem 3.24).

**Theorem 2.8.** *There exists a constant  $v > 0$  such that the following holds. Let  $f_1, \dots, f_H : \mathcal{X} \rightarrow \mathbb{R}$  be functions with  $\Delta(f_1), \dots, \Delta(f_H) \leq 1$ . For any  $\varepsilon > 0, \beta \in (0, 1)$ , and  $\tau \in \mathbb{R}$ , there exists an  $\varepsilon$ -DP algorithm such that, if  $\max_h f_h(X) \geq \tau$ , then, with probability at least  $1 - \beta$ , the algorithm outputs  $h^* \in [H]$  with the following properties:*

- $f_{h^*}(X) \geq \tau - v \cdot \log(H/\beta)/\varepsilon$ ;
- For all  $h' < h^*$ ,  $f_{h'}(X) \leq \tau + v \cdot \log(H/\beta)/\varepsilon$ .

**Exponential Mechanism.** We will also use the *exponential mechanism (EM)* of McSherry and Talwar (2007). In its generic form, EM allows us to select a solution from a candidate set  $\mathcal{H}$ . Specifically, we may define (low-sensitivity) scoring functions  $\text{scr}_h : \mathcal{X} \rightarrow \mathbb{R}$  for each  $h \in \mathcal{H}$ . Then, EM outputs a solution that approximately maximizes the score. The precise statement is given below.<sup>9</sup>

**Theorem 2.9** ((McSherry and Talwar 2007)). *For any  $\varepsilon > 0$  and  $\beta \in (0, 1)$ , a finite set  $\mathcal{H}$ , and a set of scoring functions  $\{\text{scr}_h\}_{h \in \mathcal{H}}$  such that  $\Delta(\text{scr}_h) \leq 1$  for each  $h \in \mathcal{H}$ , there is an  $\varepsilon$ -DP algorithm that, on every input  $X$ , outputs  $h^*$  such that*

$$\text{scr}_{h^*}(X) \geq \max_{h \in \mathcal{H}} \text{scr}_h(X) - \frac{2 \log(|\mathcal{H}|/\beta)}{\varepsilon}$$

*with probability at least  $1 - \beta$ .*

<sup>8</sup>The Laplace distribution with scale  $b$  is the distribution whose probability density function is proportional to  $\exp(-|x|/b)$ .

<sup>9</sup>The formulation here can be derived, e.g., by plugging  $t = \log(1/\beta)$  into Corollary 3.12 of Dwork and Roth (2014).

### 2.3 Anti-Concentration Inequality

Denote by  $\text{Ber}(1/2)$  the distribution that is 0 with probability  $1/2$ , and 1 otherwise. The following type of anti-concentration inequalities is well-known; for completeness, we provide a proof for this one in the full version of our paper (Manurangsi and Suksompong 2022b).

**Lemma 2.10.** *If  $k \geq 100$  and  $X_1, \dots, X_k$  are independent random variables drawn from  $\text{Ber}(1/2)$ , then*

$$\Pr \left[ \sum_{i=1}^k X_i < \frac{k}{2} - 0.1\sqrt{k} \right] \geq \frac{1}{4}.$$

### 3 Agent-Level DP

We begin by considering the demanding notion of agent-level DP, and provide strong negative results for this notion.

For EFC, we show a lower bound that,<sup>10</sup> when  $m > n \log n$ , holds even against  $c = \Theta\left(\sqrt{\frac{m}{n} \log n}\right)$ .

**Theorem 3.1.** *There exists a constant  $\zeta > 0$  such that, for any  $\varepsilon > 0$ , there is no agent-level  $\varepsilon$ -DP algorithm that, for any input binary additive utility functions, outputs an EFC allocation with probability higher than  $1 - \frac{\varepsilon - \zeta}{200}$ , where  $c = \left\lceil \zeta \sqrt{\frac{m}{n} \cdot \min\{\log n, \frac{m}{n}\}} \right\rceil$ .*

For proportionality, we prove a slightly weaker bound where  $c = \Theta(\sqrt{m/n})$  and the “failure probability” required for the lower bound to apply is also smaller at  $O_\varepsilon(1/n)$  (compared to  $O_\varepsilon(1)$  for envy-freeness).

**Theorem 3.2.** *There exists a constant  $\zeta > 0$  such that, for any  $\varepsilon > 0$ , there is no agent-level  $\varepsilon$ -DP algorithm that, for any input binary additive utility functions, outputs a PROPc allocation with probability higher than  $1 - \frac{\varepsilon - \zeta}{8n}$ , where  $c = \lceil \zeta \sqrt{m/n} \rceil$ .*

Due to space constraints, we only present the proof of Theorem 3.2 here. The proof of Theorem 3.1, which uses similar arguments but requires a more delicate anti-concentration inequality, can be found in the full version of our paper (Manurangsi and Suksompong 2022b).

#### 3.1 Proof of Theorem 3.2

We let  $\zeta = 0.01$ . If  $m < 100n$ , then  $c = 0$  and the theorem holds trivially even without the privacy requirement. Hence, we may assume that  $m \geq 100n$ . Throughout the proof, we consider random utility functions  $\mathbf{u} = (u_i)_{i \in N}$  where each  $u_i(j)$  is an independent  $\text{Ber}(1/2)$  random variable. For brevity, we will not repeatedly state this in the calculations below.

We start by proving the following auxiliary lemma that if  $A_i$  is small, then, for a random utility  $\mathbf{u}$  as above, the allocation fails to be PROPc for agent  $i$  with a constant probability.

**Lemma 3.3.** *For  $\zeta = 0.01$ , let  $c$  be as in Theorem 3.2 and  $A$  be an allocation such that  $|A_i| \leq m/n$ . Then, we have*

$$\Pr_{\mathbf{u}}[A \text{ is not PROPc for agent } i] \geq 1/8.$$

<sup>10</sup>Unless specified otherwise, log refers to the natural logarithm.

*Proof.* Let  $c' = 2c$ . We have

$$\begin{aligned} & \Pr_{\mathbf{u}}[A \text{ is not PROPc for agent } i] \\ & \geq \Pr_{u_i} \left[ u_i(A_i) < \frac{u_i(M)}{n} - c \right] \\ & = \Pr_{u_i} \left[ u_i(A_i) < \frac{u_i(M \setminus A_i)}{n-1} - \frac{n}{n-1} \cdot c \right] \\ & \geq \Pr_{u_i} \left[ u_i(A_i) < \frac{m}{2n} - c' \wedge u_i(M \setminus A_i) \geq \frac{m(n-1)}{2n} \right] \\ & = \Pr_{u_i} \left[ u_i(A_i) < \frac{m}{2n} - c' \right] \\ & \quad \cdot \Pr_{u_i} \left[ u_i(M \setminus A_i) \geq \frac{m(n-1)}{2n} \right] \\ & \geq \frac{1}{2} \cdot \Pr_{u_i} \left[ u_i(A_i) < \frac{m}{2n} - c' \right], \end{aligned} \tag{1}$$

where the last inequality follows from the fact that  $|M \setminus A_i| \geq m(n-1)/n$  and symmetry.

Since  $|A_i|$  is an integer,  $|A_i| \leq \lfloor m/n \rfloor$ . Moreover, since  $\lfloor m/n \rfloor \geq 100$  and the function  $f(k) = k/2 - 0.1\sqrt{k}$  is increasing in  $[1, \infty)$ , applying Lemma 2.10 with  $k = \lfloor m/n \rfloor$  gives  $\Pr_{u_i} \left[ u_i(A_i) < \frac{m}{2n} - c' \right] \geq 1/4$ . Plugging this back into (1) yields the desired bound.  $\square$

We are now ready to prove Theorem 3.2.

*Proof of Theorem 3.2.* Let  $\zeta = 0.01$  and let  $\mathcal{M}$  be any agent-level  $\varepsilon$ -DP algorithm. Consider the input utility functions  $\mathbf{u}' = (u'_i)_{i \in N}$  where the utility functions are all-zero, and consider the distribution  $\mathcal{M}(\mathbf{u}')$ . For any allocation  $A$ , we have  $\Pr_{i \in N} [|A_i| \leq m/n] \geq 1/n$  since at least one bundle  $A_i$  must have size at most  $m/n$ . This implies that  $\Pr_{i \in N, A \sim \mathcal{M}(\mathbf{u}')} [|A_i| \leq m/n] \geq 1/n$ . Thus, there exists  $i^* \in N$  such that  $\Pr_{A \sim \mathcal{M}(\mathbf{u}')} [|A_{i^*}| \leq m/n] \geq 1/n$ .

Recalling the definition of  $\mathbf{u}$  from earlier and applying Lemma 3.3, we have<sup>11</sup>

$$\begin{aligned} & \Pr_{\mathbf{u}, A \sim \mathcal{M}(\mathbf{u}')} [A \text{ is not PROPc for agent } i^*] \\ & \geq \Pr_{A \sim \mathcal{M}(\mathbf{u}')} [|A_{i^*}| \leq m/n] \\ & \quad \cdot \Pr_{\mathbf{u}, A \sim \mathcal{M}(\mathbf{u}')} [A \text{ is not PROPc for agent } i^* \mid |A_{i^*}| \leq m/n] \\ & \geq (1/n) \cdot (1/8) = 1/(8n). \end{aligned}$$

Hence, there exists  $u_{i^*}^*$  such that<sup>12</sup>

$$\Pr_{A \sim \mathcal{M}(\mathbf{u}^*)} [A \text{ is not PROPc for agent } i^* \text{ w.r.t. } u_{i^*}^*] \geq 1/(8n).$$

Now, let  $\mathbf{u}^*$  be the input utility such that  $u_{i^*}^*$  is all-zero for each  $i \neq i^*$  while  $u_{i^*}^*$  is as above. Notice that  $\mathbf{u}^*$  is adjacent to  $\mathbf{u}'$  under agent-level adjacency. Thus, applying the  $\varepsilon$ -DP guarantee of  $\mathcal{M}$ , we get

$$\Pr_{A \sim \mathcal{M}(\mathbf{u}^*)} [A \text{ is not PROPc for agent } i^* \text{ w.r.t. } u_{i^*}^*]$$

<sup>11</sup>Here, PROPc is with respect to  $\mathbf{u}$ .

<sup>12</sup>The abbreviation “w.r.t.” stands for “with respect to”.

$$\begin{aligned} &\geq e^{-\varepsilon} \cdot \Pr_{A \sim \mathcal{M}(\mathbf{u}')} [A \text{ is not PROPc for agent } i^* \text{ w.r.t. } \mathbf{u}_{i^*}^*] \\ &\geq (e^{-\varepsilon})/(8n). \end{aligned}$$

This completes the proof.  $\square$

#### 4 (Agent $\times$ Item)-Level DP

In this section, we turn our attention to (agent  $\times$  item)-level DP, which is a more relaxed notion than agent-level DP. We explore both the possibilities (Section 4.1) and limits (Section 4.2) of private algorithms with respect to this notion.

##### 4.1 Algorithms

In contrast to agent-level DP, we will show that  $O_{\varepsilon, n}(\log m)$  upper bounds can be attained in the (agent  $\times$  item)-level DP setting. Before we do so, let us first explain why straightforward approaches do not work. To this end, assume that utilities are additive and  $u_i(j) \in [0, 1]$  for all  $i \in N, j \in M$ . One may want to estimate  $u_i(S)$  for each  $S$  using the Laplace mechanism. While the Laplace mechanism guarantees that the estimate has an expected additive error of  $O(1/\varepsilon)$ , this is not useful for obtaining approximate envy-freeness or proportionality guarantees: it is possible that (almost) every good yields utility much less than 1. In this case, additive errors do not translate to any non-trivial EFC or PROPc guarantees. We will therefore develop different—and more robust—comparison methods, which ultimately allow us to overcome the aforementioned issue.

**Approximate Envy-Freeness** Our main algorithmic result for approximate envy-freeness is stated below. Note that this result holds even for non-additive utility functions.

**Theorem 4.1.** *For any  $\varepsilon > 0$  and  $\beta \in (0, 1]$ , there exists an (agent  $\times$  item)-level  $\varepsilon$ -DP algorithm that, for any input monotone utility functions, outputs a connected EFC allocation with probability at least  $1 - \beta$ , where  $c = O\left(1 + \frac{n \log(mn) + \log(1/\beta)}{\varepsilon}\right)$ .*

The high-level idea of our algorithm is to apply the exponential mechanism (Theorem 2.9) to select an allocation among the  $\leq (mn)^n$  connected allocations. This gives us an “error” in the score of  $O_\varepsilon(\log((mn)^n)) = O_\varepsilon(n \cdot \log(mn))$ . The question is how to set up the score so that (i) it has low sensitivity and (ii) such an error translates to an approximate envy-freeness guarantee. Our insight is to define the score based on the following modified utility function that “removes” a certain number of most valuable items.

**Definition 4.2.** For  $\mathbf{u} = (u_1, \dots, u_n)$  and  $k \in \mathbb{N} \cup \{0\}$ , we define  $\mathbf{u}^{-k} = (u_1^{-k}, \dots, u_n^{-k})$  by

$$u_i^{-k}(S) := \min_{T \subseteq M, |T| \leq k} u_i(S \setminus T) \quad \forall i \in N, S \subseteq M.$$

It is clear that  $\mathbf{u}^{-k}$  inherits the monotonicity of  $\mathbf{u}$ . We next list some simple but useful properties of such utility functions. The first property, whose proof is trivial, relates Definition 4.2 to approximate envy-freeness.

**Observation 4.3.** *Let  $k, d \in \mathbb{N} \cup \{0\}$ . Any allocation that is EFD with respect to  $\mathbf{u}^{-k}$  is EF( $d + k$ ) with respect to  $\mathbf{u}$ .*

The second property is that the  $d, k$  values are robust with respect to (agent  $\times$  item)-level adjacency.

**Lemma 4.4.** *Let  $k \in \mathbb{N}, d \in \mathbb{N} \cup \{0\}$ , and  $\mathbf{u}, \tilde{\mathbf{u}}$  be any two (agent  $\times$  item)-level adjacent inputs. If an allocation  $A$  is EFD with respect to  $\mathbf{u}^{-k}$ , then it is EF( $d + 2$ ) with respect to  $\tilde{\mathbf{u}}^{-(k-1)}$ .*

*Proof.* By definition of (agent  $\times$  item)-level adjacency, there exist  $i^* \in N, j^* \in M$  such that

- $u_i = \tilde{u}_i$  for all  $i \in N \setminus \{i^*\}$ , and
- $u_i(S) = \tilde{u}_i(S)$  for all  $S \subseteq M \setminus \{j^*\}$ .

Consider any  $i, i' \in N$ . Since  $A$  is EFD with respect to  $\mathbf{u}^{-k}$ , we have

$$\begin{aligned} u_i^{-k}(A_i) &\geq \min_{S \subseteq M, |S| \leq d} u_i^{-k}(A_{i'} \setminus S) \\ &= \min_{T \subseteq M, |T| \leq d+k} u_i(A_{i'} \setminus T). \end{aligned} \quad (2)$$

Furthermore, we have

$$\begin{aligned} u_i^{-k}(A_i) &= \min_{T \subseteq M, |T| \leq k} u_i(A_i \setminus T) \\ &\leq \min_{T \subseteq M, |T| \leq k-1} u_i(A_i \setminus (T \cup \{j^*\})) \\ &= \min_{T \subseteq M, |T| \leq k-1} \tilde{u}_i(A_i \setminus (T \cup \{j^*\})) \\ &\leq \min_{T \subseteq M, |T| \leq k-1} \tilde{u}_i(A_i \setminus T) = \tilde{u}_i^{-(k-1)}(A_i). \end{aligned} \quad (3)$$

Moreover,

$$\begin{aligned} &\min_{T \subseteq M, |T| \leq d+k} u_i(A_{i'} \setminus T) \\ &\geq \min_{T \subseteq M, |T| \leq d+k} u_i(A_{i'} \setminus (T \cup \{j^*\})) \\ &= \min_{T \subseteq M, |T| \leq d+k} \tilde{u}_i(A_{i'} \setminus (T \cup \{j^*\})) \\ &\geq \min_{T \subseteq M, |T| \leq d+k+1} \tilde{u}_i(A_{i'} \setminus T) \\ &= \min_{S \subseteq M, |S| \leq d+2} \tilde{u}_i^{-(k-1)}(A_{i'} \setminus S). \end{aligned} \quad (4)$$

Thus, we can conclude that

$$\begin{aligned} \tilde{u}_i^{-(k-1)}(A_i) &\stackrel{(3)}{\geq} u_i^{-k}(A_i) \\ &\stackrel{(2)}{\geq} \min_{T \subseteq M, |T| \leq d+k} u_i(A_{i'} \setminus T) \\ &\stackrel{(4)}{\geq} \min_{S \subseteq M, |S| \leq d+2} \tilde{u}_i^{-(k-1)}(A_{i'} \setminus S). \end{aligned}$$

It follows that  $A$  is EF( $d + 2$ ) with respect to  $\tilde{\mathbf{u}}^{-(k-1)}$ .  $\square$

We can now define our scoring function.

**Definition 4.5.** For an allocation  $A, \mathbf{u} = (u_1, \dots, u_n)$ , and  $g \in \mathbb{N}$ , define

$$\text{scr}_A^g(\mathbf{u}) := -\min\{t \in [g] \mid A \text{ is EF}(2t) \text{ w.r.t. } \mathbf{u}^{-(g-t)}\}.$$

We let  $\text{scr}_A^g(\mathbf{u}) = -g$  if the set above is empty.

---

Algorithm 1: (Agent  $\times$  item)-level  $\varepsilon$ -DP algorithm for EF $c$

---

**Parameter:**  $\varepsilon > 0, \beta \in (0, 1]$

- 1:  $g \leftarrow 4 \left\lceil 1 + \frac{\log((mn)^n/\beta)}{\varepsilon} \right\rceil$
  - 2: **return** the allocation output by the  $\varepsilon$ -DP exponential mechanism using the scoring function  $\text{scr}_A^g$  (Definition 4.5) with the candidate set  $\mathcal{P}^{\text{conn}}(m, n)$  (Section 2.1).
- 

The following lemma shows that this scoring function has low sensitivity.

**Lemma 4.6.** *For any allocation  $A$  and  $g \in \mathbb{N}$ ,  $\Delta(\text{scr}_A^g) \leq 1$ .*

*Proof.* Let  $\mathbf{u}, \tilde{\mathbf{u}}$  be any pair of (agent  $\times$  item)-level adjacent inputs. Assume without loss of generality that  $\text{scr}_A^g(\mathbf{u}) \geq \text{scr}_A^g(\tilde{\mathbf{u}})$ . Let  $t^* := -\text{scr}_A^g(\mathbf{u})$ .

If  $t^* = g$ , then  $\text{scr}_A^g(\tilde{\mathbf{u}}) \leq \text{scr}_A^g(\mathbf{u}) = -g$ , so  $\text{scr}_A^g(\tilde{\mathbf{u}}) = -g$ . Otherwise,  $t^* \leq g - 1$ , and  $A$  is EF( $2t^*$ ) with respect to  $\mathbf{u}^{-(g-t^*)}$ . Lemma 4.4 ensures that  $A$  is EF( $2(t^* + 1)$ ) with respect to  $\tilde{\mathbf{u}}^{-(g-t^*-1)}$ . Thus,  $\text{scr}_A^g(\tilde{\mathbf{u}}) \geq -(t^* + 1)$ . Since  $-t^* = \text{scr}_A^g(\mathbf{u}) \geq \text{scr}_A^g(\tilde{\mathbf{u}})$ , we have  $|\text{scr}_A^g(\mathbf{u}) - \text{scr}_A^g(\tilde{\mathbf{u}})| \leq 1$  in both cases, completing the proof.  $\square$

With all the ingredients ready, we can prove Theorem 4.1 by applying the exponential mechanism with appropriate parameters (see Algorithm 1).

*Proof of Theorem 4.1.* Let  $g = 4 \left\lceil 1 + \frac{\log((mn)^n/\beta)}{\varepsilon} \right\rceil$ . We run the exponential mechanism using the scoring function  $\text{scr}_A^g$  with the candidate set  $\mathcal{P}^{\text{conn}}(m, n)$ . By Theorem 2.9, this is an  $\varepsilon$ -DP algorithm that, for each  $\mathbf{u}$ , with probability at least  $1 - \beta$ , outputs an allocation  $A^*$  such that

$$\text{scr}_{A^*}^g(\mathbf{u}) \geq \max_{A \in \mathcal{P}^{\text{conn}}(m, n)} \text{scr}_A^g(\mathbf{u}) - \frac{2 \log \left( \frac{|\mathcal{P}^{\text{conn}}(m, n)|}{\beta} \right)}{\varepsilon}.$$

Fix any  $\mathbf{u}$ , and define  $A^*$  as above. By Theorem 2.1, there exists a connected allocation  $A^{\text{EF2}}$  that is EF2 with respect to  $\mathbf{u}^{-(g-1)}$ . This means that  $\text{scr}_{A^{\text{EF2}}}^g = -1$ . Furthermore, we have<sup>13</sup>  $|\mathcal{P}^{\text{conn}}(m, n)| \leq (mn)^n$ . Plugging these into the inequality above, we get

$$\text{scr}_{A^*}^g(\mathbf{u}) \geq -1 - \frac{2 \log((mn)^n/\beta)}{\varepsilon} \geq -\frac{g}{2},$$

where the latter inequality follows from our choice of  $g$ . Hence,  $A^*$  is EF( $2c$ ) with respect to  $\mathbf{u}^{-(g-c)}$  for some  $c \leq g/2$ . Invoking Observation 4.3, we find that  $A^*$  is EF( $g + c$ ), and therefore EF( $3g/2$ ), with respect to  $\mathbf{u}$ . This concludes our proof.  $\square$

<sup>13</sup>Indeed, from the set  $M = \{1, 2, \dots, m\}$ , we can allocate one “block” of items at a time starting from items with lower indices. There are at most  $m$  possibilities for the size of the next block, this block can be allocated to one of the (at most  $n$ ) remaining agents, and we allocate  $n$  blocks in total, hence the bound  $(mn)^n$ .

**Approximate Proportionality** Next, we present an improved result for approximate proportionality, where the dependence on  $n$  is reduced to  $O(\log n)$ .

**Theorem 4.7.** *For any  $\varepsilon > 0$  and  $\beta \in (0, 1]$ , there exists an (agent  $\times$  item)-level  $\varepsilon$ -DP algorithm that, for any input additive utility functions, outputs a connected PROP $c$  allocation with probability at least  $1 - \beta$ , where  $c = O\left(\log n + \frac{\log(mn/\beta)}{\varepsilon}\right)$ .*

Our algorithm is based on the well-known “moving-knife” procedure from cake cutting (Dubins and Spanier 1961). A natural way to implement this idea in our setting is to place the items on a line, put a knife at the left end, and move it rightwards until some agent values the subset of items to the left of the knife at least  $1/n$  of her value for the whole set of items. We give this subset to this agent, and proceed similarly with the remaining agents and items. To make this procedure DP, we can replace the check of whether each agent receives sufficiently high value with the SVT algorithm (Theorem 2.8), where the usual utility is modified similarly to Definition 4.5 to achieve low sensitivity. While this approach is feasible, it does not establish the bound we want: since the last agent has to participate in  $n$  “rounds” of this protocol, the basic composition theorem (Theorem 2.5) implies that we can only allot a privacy budget of  $\varepsilon/n$  in each round. This results in a guarantee of the form  $c = O(\log m/(\varepsilon/n)) = O((n \log m)/\varepsilon)$ , which does not distinctly improve upon the guarantee in Theorem 4.1.

To overcome this issue, notice that instead of targeting a single agent, we can continue moving our knife until at least  $n/2$  agents value the subset of items to the left of the knife at least half of the entire set.<sup>14</sup> This allows us to recurse on both sides, thereby reducing the number of rounds to  $\log n$ . Hence, we may allot a privacy budget of  $\varepsilon/\log n$  in each round. Unfortunately, this only results in a bound of the form  $c = O(\log m/(\varepsilon/\log n)) = O((\log n \log m)/\varepsilon)$ , which is still worse than what we claim in Theorem 4.7.

Our last observation is that we can afford to make more mistakes in earlier rounds: for example, in the first round, we would be fine with making an “error” of roughly  $O(n)$  in the knife position because the subsets on both sides will be subdivided to  $\Omega(n)$  parts later. As a result, our strategy is to allot less privacy budget in earlier rounds and more in later rounds. By letting the privacy budgets form a geometric sequence, we can achieve our claimed  $O(\log(mn)/\varepsilon)$  bound. The detailed proof can be found in the full version of our paper (Manurangsi and Suksompong 2022b).

## 4.2 Lower Bounds

Next, we prove lower bounds for (agent  $\times$  item)-level DP via the *packing method* (Hardt and Talwar 2010). This involves constructing inputs that are close to one another (with respect to the corresponding adjacency notion) such that the acceptable solutions (i.e., EF $c$  or PROP $c$  allocations) are different for different inputs. The DP requirement can then be used to rule out the existence of algorithms with strong utility guarantees. We reiterate that our lower bounds hold only

<sup>14</sup>Even and Paz (1984) used a similar idea in cake cutting.

against connected allocations. In our constructions, we design the utility functions so that each input forces us to pick particular positions to cut in order to get an EFC or PROPC allocation. We start with the proof for envy-freeness.

**Theorem 4.8.** *There exists  $\zeta > 0$  such that, for any  $\varepsilon \in (0, 1]$ , there is no  $\varepsilon$ -DP algorithm that, for any input binary additive utility functions, outputs a connected EFC allocation with probability at least 0.5, where  $c = \left\lceil \zeta \cdot \min \left\{ \frac{\log m}{\varepsilon}, \frac{m}{n}, \sqrt{m} \right\} \right\rceil$ .*

*Proof.* Let  $\zeta = 0.01$ ,  $c$  be as in the theorem statement, and  $T = \lfloor m/(4c + 4) \rfloor$ . We may assume that  $c \geq 1$ , as otherwise the theorem holds trivially even without the privacy requirement. Consider the following utility functions.

- Let  $\mathbf{u}' = (u'_1, \dots, u'_n)$  denote the binary additive utility functions defined as follows:
  - $u'_1$  and  $u'_2$  are all-zero utility functions.
  - For all  $i \in \{3, \dots, n\}$  and  $j \in M$ , let

$$u'_i(j) = \begin{cases} 1 & \text{if } j \geq m - (c + 1)(n - 2); \\ 0 & \text{otherwise.} \end{cases}$$

- For every  $t \in [T]$ , let  $\mathbf{u}^t = (u^t_1, \dots, u^t_n)$  denote the binary additive utility functions defined as follows:
  - $u^t_1$  and  $u^t_2$  are defined as follows:

$$u^t_1(j) = u^t_2(j) = \begin{cases} 1 & \text{if } \left\lfloor \frac{j-1}{2c+1} \right\rfloor = t - 1; \\ 0 & \text{otherwise,} \end{cases}$$

for all  $j \in M$ .

- For all  $i \in \{3, \dots, n\}$ ,  $u^t_i$  is exactly the same as  $u'_i$  defined earlier.

Suppose for contradiction that there is an  $\varepsilon$ -DP algorithm  $\mathcal{M}$  that, with probability at least 0.5, outputs a connected allocation that is EFC for its input utility functions. For each  $t \in [T]$ , let  $\mathcal{A}^t$  denote the set of allocations that are EFC for  $\mathbf{u}^t$ . The assumption on  $\mathcal{M}$  can be written as

$$\Pr[\mathcal{M}(\mathbf{u}^t) \in \mathcal{A}^t] \geq 0.5. \quad (5)$$

Let  $\sim$  denote the (agent  $\times$  item)-level adjacency relation. One can check that  $\mathbf{u}' \sim_{4c+2} \mathbf{u}^t$  for all  $t \in [T]$ . Using this fact together with group DP (Lemma 2.6), we have

$$\begin{aligned} \Pr[\mathcal{M}(\mathbf{u}') \in \mathcal{A}^t] &\geq e^{-\varepsilon(4c+2)} \cdot \Pr[\mathcal{M}(\mathbf{u}^t) \in \mathcal{A}^t] \\ &\stackrel{(5)}{\geq} 0.5 \cdot e^{-\varepsilon(4c+2)}. \end{aligned} \quad (6)$$

**Lemma 4.9.**  $\mathcal{A}^1, \dots, \mathcal{A}^T$  are disjoint.

Due to space constraints, we defer the proof of Lemma 4.9 to the full version of our paper (Manurangsi and Suksompong 2022b). The idea behind the proof is that, if there were an allocation  $A$  contained in both  $\mathcal{A}^t$  and  $\mathcal{A}^{t'}$  for some  $t < t'$ , there would have to be at least  $n - 3$  cuts between item  $m - (c + 1)(n - 2)$  and item  $m$ , a cut between item  $(2c + 1)(t - 1) + 1$  and item  $(2c + 1)t$ , and a cut between item  $(2c + 1)(t' - 1) + 1$  and item  $(2c + 1)t'$ . Since these

intervals of items are pairwise disjoint, the  $n - 1$  cuts must be all the cuts in  $A$ . Based on this, we can then show that no matter which agent the leftmost interval of  $A$  is assigned to,  $A$  cannot be EFC for both  $\mathbf{u}^t$  and  $\mathbf{u}^{t'}$  simultaneously.

Lemma 4.9 implies that

$$\begin{aligned} 1 &\geq \sum_{t \in [T]} \Pr[\mathcal{M}(\mathbf{u}') \in \mathcal{A}^t] \\ &\stackrel{(6)}{\geq} 0.5 \cdot e^{-\varepsilon(4c+2)} \cdot T \\ &\geq 0.5 \cdot e^{-6c\varepsilon} \cdot \left\lfloor \frac{m}{8c} \right\rfloor \\ &\geq 0.5 \cdot e^{-0.06 \log m} \cdot \lfloor 10\sqrt{m} \rfloor \\ &\geq (0.5/\sqrt{m}) \cdot (5\sqrt{m}) > 1, \end{aligned}$$

where the third and fourth inequalities follow from our choice of parameters  $c, T$  and the assumption that  $c \geq 1$ . This is a contradiction which establishes Theorem 4.8.  $\square$

For proportionality, we obtain a slightly weaker bound where the  $\log m$  term is replaced by  $\log(m/n)/n$ . The proof is similar to that of Theorem 4.8 and can be found in the full version of our paper (Manurangsi and Suksompong 2022b).

**Theorem 4.10.** *There exists a constant  $\zeta > 0$  such that, for any  $\varepsilon \in (0, 1]$ , there is no  $\varepsilon$ -DP algorithm that, for any input binary additive utility functions, outputs a connected PROPC allocation with probability at least 0.5, where  $c = \left\lceil \zeta \cdot \min \left\{ \frac{\log(m/n)}{\varepsilon n}, \frac{m}{n}, \sqrt{\frac{m}{n}} \right\} \right\rceil$ .*

## 5 Conclusion and Future Work

In this paper, we have studied the fundamental task of fair division under differential privacy constraints, and provided algorithms and impossibility results for approximate envy-freeness and proportionality. There are several open questions left by our work. First, it would be useful to close the gaps in terms of  $n$ ; for example, our envy-freeness upper bound for (agent  $\times$  item)-level DP grows linearly in  $n$  (Theorem 4.1) but our lower bound (Theorem 4.8) does not exhibit this behavior. Another perhaps more interesting technical direction is to extend our lower bounds for (agent  $\times$  item)-level DP to arbitrary (i.e., not necessarily connected) allocations. Specifically, we leave the following intriguing open question: Is there an (agent  $\times$  item)-level  $\varepsilon$ -DP algorithm that, with probability at least 0.99, outputs an EFC allocation for  $c = O_\varepsilon(1)$  regardless of the values of  $n$  and  $m$ ?

While we have considered the original notion of DP proposed by Dwork et al. (2006b), there are a number of modifications that could be investigated in future work. A commonly studied notion is *approximate DP* (also called  $(\varepsilon, \delta)$ -DP), which has an additional parameter  $\delta \geq 0$  that specifies the probability with which the condition  $\Pr[\mathcal{M}(X) = o] \leq e^\varepsilon \cdot \Pr[\mathcal{M}(X') = o]$  is allowed to fail (Dwork et al. 2006a). The notion of DP that we use in this paper corresponds to the case  $\delta = 0$  and is often referred to as *pure DP*. Several problems in the literature are known to admit approximate-DP algorithms with better guarantees compared to pure-DP algorithms (see, e.g., the work of Steinke and Ullman (2016)). In light of this, it would be interesting to explore whether a similar phenomenon occurs in fair division as well.

## Acknowledgments

This work was partially supported by the Singapore Ministry of Education under grant number MOE-T2EP20221-0001 and by an NUS Start-up Grant.

## References

- Abowd, J. M. 2018. The US Census Bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, 2867.
- Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; and Voudouris, A. A. 2022. Fair division of indivisible goods: a survey. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 5385–5393.
- Apple Differential Privacy Team. 2017. Learning with privacy at scale. <http://machinelearning.apple.com/research/learning-with-privacy-at-scale>. Accessed 2022-11-21.
- Aziz, H.; Li, B.; Moulin, H.; and Wu, X. 2022. Algorithmic fair allocation of indivisible items: a survey and new questions. *ACM SIGecom Exchanges*, 20(1): 24–40.
- Bei, X.; Igarashi, A.; Lu, X.; and Suksompong, W. 2022. The price of connectivity in fair division. *SIAM Journal on Discrete Mathematics*, 36(2): 1156–1186.
- Bilò, V.; Caragiannis, I.; Flammini, M.; Igarashi, A.; Monaco, G.; Peters, D.; Vinci, C.; and Zwicker, W. S. 2022. Almost envy-free allocations with connected bundles. *Games and Economic Behavior*, 131: 197–221.
- Bouveret, S.; Cechlárová, K.; Elkind, E.; Igarashi, A.; and Peters, D. 2017. Fair division of a graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 135–141.
- Bouveret, S.; Chevaleyre, Y.; and Maudet, N. 2016. Fair allocation of indivisible goods. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 12, 284–310. Cambridge University Press.
- Brams, S. J.; and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Ding, B.; Kulkarni, J.; and Yekhanin, S. 2017. Collecting telemetry data privately. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NIPS)*, 3571–3580.
- Dubins, L. E.; and Spanier, E. H. 1961. How to cut a cake fairly. *American Mathematical Monthly*, 68(1): 1–17.
- Dwork, C. 2008. Differential privacy: a survey of results. In *Proceedings of the 5th International Conference on Theory and Applications of Models of Computation (TAMC)*, 1–19.
- Dwork, C.; Kenthapadi, K.; McSherry, F.; Mironov, I.; and Naor, M. 2006a. Our data, ourselves: Privacy via distributed noise generation. In *Proceedings of the 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, 486–503.
- Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. D. 2006b. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Theory of Cryptography Conference (TCC)*, 265–284.
- Dwork, C.; Naor, M.; Reingold, O.; Rothblum, G. N.; and Vadhan, S. P. 2009. On the complexity of differentially private data release: Efficient algorithms and hardness results. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, 381–390.
- Dwork, C.; and Roth, A. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4): 211–407.
- Erlingsson, Ú.; Pihur, V.; and Korolova, A. 2014. Rap-  
por: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 1054–1067.
- Even, S.; and Paz, A. 1984. A note on cake cutting. *Discrete Applied Mathematics*, 7(3): 285–296.
- Greenberg, A. 2016. Apple’s ‘differential privacy’ is about collecting your data—but not *your* data. <http://www.wired.com/2016/06/apples-differential-privacy-collecting-data>. Accessed 2022-08-08.
- Hardt, M.; and Talwar, K. 2010. On the geometry of differential privacy. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, 705–714.
- Igarashi, A. 2023. How to cut a discrete cake fairly. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*. Forthcoming.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Economics and Computation (EC)*, 125–131.
- Manurangsi, P.; and Suksompong, W. 2022a. Almost envy-freeness for groups: Improved bounds via discrepancy theory. *Theoretical Computer Science*, 930: 179–195.
- Manurangsi, P.; and Suksompong, W. 2022b. Differentially private fair division. *CoRR*, abs/2211.12738.
- Markakis, E. 2017. Approximation algorithms and hardness results for fair division with indivisible goods. In Endriss, U., ed., *Trends in Computational Social Choice*, chapter 12, 231–247. AI Access.
- McSherry, F.; and Talwar, K. 2007. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 94–103.
- Moulin, H. 2003. *Fair Division and Collective Welfare*. MIT Press.
- Moulin, H. 2019. Fair division in the Internet age. *Annual Review of Economics*, 11: 407–441.
- Shankland, S. 2014. How Google tricks itself to protect Chrome user privacy. <http://www.cnet.com/tech/services-and-software/how-google-tricks-itself-to-protect-chrome-user-privacy>. Accessed 2022-11-21.
- Steinke, T.; and Ullman, J. R. 2016. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7(2): 3–22.
- Suksompong, W. 2019. Fairly allocating contiguous blocks of indivisible items. *Discrete Applied Mathematics*, 260: 227–236.



- Suksompong, W. 2021. Constraints in fair division. *ACM SIGecom Exchanges*, 19(2): 46–61.
- Sun, N.; and Yang, Z. 2009. Strategy proof and privacy preserving fair allocation mechanism. *Japanese Economic Review*, 60(2): 143–151.
- Vadhan, S. 2017. The complexity of differential privacy. In Lindell, Y., ed., *Tutorials on the Foundations of Cryptography*, chapter 7, 347–450. Springer International Publishing.
- Walsh, T. 2020. Fair division: the computer scientist’s perspective. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 4966–4972.