Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Truthful Cake Cutting Mechanisms with Externalities: Do Not Make Them Care for Others Too Much!

Minming Li
City University of Hong Kong
Hong Kong, China
minming.li@cityu.edu.hk

Jialin Zhang
Institute of Computing Technology
Beijing, China
zhangjl2002@gmail.com

Qiang Zhang
University of Warsaw
Warsaw, Poland
qzhang@mimuw.edu.pl

Abstract

We study truthful mechanisms in the context of cake cutting when agents not only value their own pieces of cake but also care for the pieces assigned to other agents. In particular, agents derive benefits or costs from the pieces of cake assigned to other agents. This phenomenon is often referred to as positive or negative externalities. We propose and study the following model: given an allocation, externalities of agents are modeled as percentages of the reported values that other agents have for their pieces. We show that even in this restricted class of externalities, under some natural assumptions, no truthful cake cutting mechanisms exist when externalities are either positive or negative. However, when the percentages agents get from each other are small, we show that there exists a truthful cake cutting mechanism with other desired properties.

1 Introduction

Cake cutting, a fundamental resource allocation problem, has been extensively studied by economists and mathematicians in terms of fairness from the last century. Among the literature of cake cutting, *envy-freeness* and *proportionality* are two most notable concepts to capture fairness. With respect to an allocation, a mapping between pieces of cake and agents, envy-freeness means agents weakly prefer their own pieces to the ones assigned to other agents, and proportionality implies that each agent among n agents has a piece of cake that is at least $1/n$ of his value for the entire cake. With the increasing attention across artificial intelligence and algorithmic mechanism design [Nisan and Ronen, 1999], cake cutting has been considered in multi-agent systems recently, particularly when agents have private values for different pieces of cake. For instance, Alice prefers the pieces with strawberries while Bob prefers the ones with chocolates. Based on values declared by agents, a cake is cut into pieces, which are then assigned to agents. In most of the cases, we would like the allocation to satisfy some good properties such as envy-freeness and/or proportionality. However, it may motivate selfish agents to benefit by misreporting their true values. A cake cutting mechanism is *truthful* if it prevents agents from benefiting by reporting false values in any circumstance. When agents' values are formalized as so-called valuation functions, one could expect that it is difficult to have truthful cake cutting mechanisms with good properties if the domain of agents' valuation functions is unrestricted. Existing works [Chen *et al.*, 2013] showed a truthful deterministic (resp. randomized) mechanism exists if agents' valuation functions are *piecewise uniform* (resp. *piecewise linear*). In this paper, we restrict our study within the domain of piecewise uniform valuation functions.

This paper takes one step forward by studying the design of truthful cake cutting mechanisms when externalities are considered. In economics, externalities are defined as the benefits or costs that affect a party which was not chosen to incur [Baumol, 1972]. While externalities may often be neglected, studies have shown their importance in a wide range of settings. For example, revenues in auctions [Haghpanah *et al.*, 2013], or supply chain inventory management [Cachon, 1999; Netessine and Zhang, 2005]. In this paper, when externalities exist in the context of cake cutting, we consider that agents derive benefits or costs from the pieces assigned to others. Externalities make the values of agents for an allocation depend on the social assignment rather than their individual assignments.

1.1 Model and Truthfulness

There are different approaches to modeling externalities in different settings. To the best of our knowledge, we are the first to model externalities in the setting of cake cutting as follows. Agents could derive benefits or costs from each other, which are certain percentages of the *reported values* of other agents for their own pieces. That is, given an allocation where each agent gets a piece of cake, the value an agent has for the allocation is his valuation for his piece of cake plus (or minus when externalities are negative) certain percentages of other agents' valuations for their pieces of cake. Although this model is somehow restricted, it does capture some real-life applications. For example, revenue-sharing contracts are common in supply chains. Under a revenue-sharing contract, a retailer pays a supplier a wholesale price, plus a percentage of the revenue the retailer has. For more details on revenue-sharing contracts we refer the reader to [Cachon and Lariviere, 2005]. Negative externalities happen when two rivals in a market are competing for advertisement slots. One company increases its market share by gaining advertisement slots but it also loses some market share if advertisement slots are assigned to its rival. One could model the market share lost by the company as a certain percentage of the market share gain by its rivals. Note that this allows that the valuation of an agent for an allocation depends on the valuations reported by other agents. It creates a situation in which agents do not fully know their own valuations for allocations. However, since we are interested in truthful mechanisms in the sense of truth-telling being a dominant strategy for all agents, it guarantees that agents do not have incentive to deviate their reports after seeing the reports by other agents. Another approach closely related to our model is to model the externalities as certain percents of the *true values* of other agents for their own pieces. In this paper, we choose the former approach due to the following two reasons. First, in the latter approach, the valuation of an agent is affected both by the reported valuations of other agents (which affect the allocation returned by the mechanism) and by the true valuations of other agents (which affect the externalities he obtains). To obtain the valuations of allocations, agents need to know the private information (true valuations) of other agents, which may be impossible. Second, once truthful mechanisms exist in the former approach, the results also imply that truth-telling is a Nash equilibrium in the latter approach, which is a reasonable solution concept for the latter approach.

1.2 Assumptions

Before we continue, we would like to take a tour of the assumptions and properties that mechanisms are assumed to satisfy in this paper.

In classical settings without externalities, normalizations are often carried out since it does not affect the preferences of agents, that is, agents always prefer a larger piece of cake in some sense. However, when externalities exist, agents care for their own assignments but also for the assignments of other agents. It is not precise to say that agents always want to maximize their pieces of cake. Hence, in this paper, we only normalize the valuations of agents for the entire cake to 1. In this way, each component in the valuation of an agent for an allocation is normalized (see Section 2.1 for the concrete model). Hence, in general we could say agents want to maximize the pieces of cake (i.e., the assignments) that are important to them. This matches well with the settings without externalities in which the only important pieces of cake for agents are the ones assigned to them.

For mechanisms, besides truthfulness, envy-freeness and proportionality mentioned above, we are also interested in mechanisms that are *non-wasteful* and *position independent*. All of these properties or similar properties could be found in the literature, but some of them are not well defined in cake cutting settings when externalities exist. A mechanism is non-wasteful if it assigns any piece of cake, which is desired by at least one agent, to some agent who desires it. By relaxing it, a trivial mechanism could be one that assigns nothing to all agents. Position independence can be seen as a similar concept to anonymity to avoid dictatorship mechanisms. Anonymity is ambiguous in two ways when externalities exist. Generally, anonymity means that the outcome of a mechanism does not depend on the identities of agents. However, when externalities exist, it is necessary to know the identities of other agents in order to evaluate the externalities obtained by an agent. The second ambiguity is that outcomes of mechanisms could be defined in terms of either assignments or values of agents when externalities are introduced. In this paper, we say a mechanism is *position independent* if agents' values over the allocation get permuted accordingly when agents are permuted in the input to the mechanism. One could easily see that this definition is more general than the definition in terms of assignments. Permuted allocation ensures the values of agents over the allocation remain the same, while, vice versa it may not be necessarily true. Additionally, randomized mechanisms are necessary if one defines position independence in terms of assignments. Considering two agents having the same valuation for the cake, randomization is required since any deterministic mechanism cannot distinguish the identities of the agents.

1.3 Our Results

In this paper, we study truthful cake cutting mechanisms when externalities are modeled in the way that agents are entitled benefits or costs from each other's assignment. Our main contributions are two-fold.

- First, we show that, in general no truthful, non-wasteful and position independent cake cutting mechanisms exist when externalities are either positive or negative.
- On the other hand, for positive externalities, when agents have limited benefits from others, we show that there exists a truthful cake cutting mechanism with other desired properties.

1.4 Related Work

A detailed review on cake cutting from the perspectives of economics and mathematics refers to Robertson and Webb [Robertson and Webb, 1998]. Besides, there is a recent survey [Procaccia, 2013] that summarizes cake cutting algorithms from both computational and game-theoretic viewpoints. In addition, there are series of papers studying cake cutting algorithms from different aspects in computer science and AI communities, e.g. fairness [Balkanski *et al.*, 2014; Brânzei *et al.*, 2013; Caragiannis *et al.*, 2011; Kurokawa *et al.*, 2013; Segal-Halevi *et al.*, 2015], social welfare [Bei *et al.*, 2012; Cohler *et al.*, 2011], equilibrium [Brânzei and Miltersen, 2013], truthfulness [Chen *et al.*, 2013]. However, none of these papers consider truthful cake cutting mechanisms when externalities exist.

Two most relevant papers to us are [Chen *et al.*, 2013] and [Brânzei *et al.*, 2013]. Chen *et al.* [2013] studied truthful cake cutting mechanisms without considering externalities. They gave a deterministic truthful, proportional, envy-free cake cutting mechanism for piecewise uniform valuation functions, and a randomized truthful-in-expectation, proportional, envy-free cake cutting mechanism for piecewise linear valuation functions. On the other hand, Brânzei *et al.* [2013] studied the concepts of envy-freeness when externalities exist in the context of cake cutting but they did not explicitly describe the externalities. Their main contributions are the following two concepts: *swap envy-freeness* and *swap stability*.

In a swap envy-free outcome, an agent cannot benefit from swapping his assignment with that of of another agent. In a swap stable outcome, an agent cannot benefit from swapping the assignments of any pair of agents. We adopt the swap-envy-freeness to generalize the classical definition of envy-freeness in cake cutting settings when externalities exist.

2 Preliminaries

2.1 Model

The entire cake is the interval $[0, 1]$. A piece of cake is a set of disjoint subintervals in $[0, 1]$. There is a set $N = \{1, \ldots, n\}$ of agents who have values for different pieces of cake. Each agent $i \in N$ has a private non-negative valuation density function $v_i(\cdot)$ over the entire cake. We say an agent *prefers* a piece of cake if his valuation density is positive across the entire piece. The value of agent i for a piece x of cake is given by $V_i(x) = \sum_{I \in x} \int_I v_i(z)dz$. The definition allows us to consider pieces of cake intersecting at the boundaries as disjoint pieces since agents have no values for any single point of cake. For example, piece $([\frac{1}{2}, \frac{3}{5}], [\frac{7}{8}, 1])$ and piece $[\frac{3}{5}, \frac{7}{8}]$ are disjoint. Similar to classical settings, the values of agents for the entire cake are normalized, that is, $V_i([0, 1]) = 1$ for all $i \in N$. In this paper, we focus on *piecewise uniform valuation functions*, where each agent prefers a piece of cake and has the same marginal value over this piece. Formally, V_i is piecewise uniform if and only if v_i is either some constant $c \in \mathbb{R}^+$ or zero across the entire cake. We denote $\mathcal{V}$ as the class of piecewise uniform valuation functions. In this paper, due to the property of piecewise uniform valuation functions and $V_i([0, 1]) = 1$, we say agent i declares his preferred pieces (or equivalently, intervals) of cake denoted by I_i instead of declaring the valuation density function v_i. Let $|I_i|$ be the total length of the intervals agent i prefers. When S is a set of agents, $|I_S|$ is the total length of the intervals which at least one agent in S prefers.

An allocation $A = (A_1, \ldots, A_n)$ that consists of n pieces of cake is feasible if and only if all pieces $A_1, \ldots, A_n$ are disjoint. In allocation A, agent i gets the piece A_i of cake. We model the externalities by that agents derive benefits or costs that are some percent of the values other agents report. Specifically, agent i derives a value $\alpha_{i,j} V_j(A_j)$ when A_j is allocated to agent j. Note that $\alpha_{i,j} \in \mathbb{R}$. With externalities, the value of agent i for allocation A is $V_i(A) = V_i(A_i) + \sum_{j \neq i}^{N} \alpha_{i,j} V_j(A_j) = \sum_{j \in N} \alpha_{i,j} V_j(A_j)$ where $\alpha_{i,i} = 1$. Note that $V_i(A)$ is not normalized in this paper as discussed in the introduction.

For convenience, we use $V_i(\cdot)$ to denote both the value of agent i for a piece of cake and the value of agent i for an allocation. When z is a piece of cake, $V_i(z)$ is the value of agent i for the piece z of cake. When z is an allocation, $V_i(z)$ is the value of agent i for allocation z.

2.2 Properties of Cake Cutting Algorithms

A (deterministic) cake cutting mechanism M maps the entire cake and valuation functions to a feasible allocation A, and the rest of cake $[0, 1] \backslash \bigcup_{i \in N} A_i$ is disposed at zero cost. Given valuations $V_1, \ldots, V_n$, we denote $M(V_1, \ldots, V_n)$ as the allocation produced by M, and $M_i(V_1, \ldots, V_n)$ as the piece of cake assigned to agent i.

A cake cutting mechanism M is *truthful* if for every $i \in N$, every $V_1, \ldots, V_n \in \mathcal{V}$ and every $V_i' \in \mathcal{V}$, it holds that $V_i(M(V_1, \ldots, V_n)) \geq V_i(M(V_1, \ldots, V_{i-1}, V_i', V_{i+1}, V_n))$. The truthfulness guarantees that agents cannot benefit if they misreport their true valuation functions regardless of the reports from other agents. Note that agent i's valuation function V_i also contains all $\alpha_{i,j}$ for $j \neq i$. One can interpret $\{\alpha_{i,j}\}$ as public or private information. In this paper, we assume they are public knowledge. It would not be difficult to see that there is little we can achieve if $\{\alpha_{i,j}\}$ is unrestricted and private. Since the valuation function is piecewise uniform, the only private information of agent i is the piece of cake I_i agent i prefers.

A cake cutting mechanism M is *proportional* if for every $i \in N$, every $V_1, \ldots, V_n \in \mathcal{V}$, it holds that $V_i(M(V_1, \ldots, V_n)) \geq \frac{1}{n} V_i(\tilde{A}_i)$, where $\tilde{A}_i$ is the allocation that maximizes agent i's value. Due to externalities, $\tilde{A}_i$ may not be equal to giving all agent i's preferred intervals to him.

A cake cutting mechanism M is *swap envy-free* if for every $i, j \in N$, every $V_1, \ldots, V_n \in \mathcal{V}$, it holds that $V_i(M_i(V_1, \ldots, V_n)) + \alpha_{i,j} V_j(M_j(V_1, \ldots, V_n)) \geq V_i(M_j(V_1, \ldots, V_n)) + \alpha_{i,j} V_j(M_i(V_1, \ldots, V_n))$.

A cake cutting mechanism M is *non-wasteful* if it always allocates every piece of cake that is preferred by at least one agent to some agent who prefers it.

A cake cutting mechanism M is *position independent* if for every $V_1, \ldots, V_n \in \mathcal{V}$ and every permutation π of agents, it holds for any $i \in N$, $V_i(M(V_1, \ldots, V_n)) = \tilde{V}_{\pi(i)}(M(\tilde{V}_1, \ldots, \tilde{V}_n)$ where $\tilde{V}_j = V_{\pi(j)}$.

3 Impossibility Result

We begin with the following impossibility result when externalities are negative.

Theorem 1. *No truthful, non-wasteful and position independent cake cutting mechanisms exist when $\{\alpha_{i,j}\}$ are negative.*

Proof. Suppose there exists a truthful, non-wasteful and position independent cake cutting mechanism M. Assume that there are two agents and $\alpha_{1,2}, \alpha_{2,1} < 0$. Let $x, y, z \subset [0, 1]$ be three nonempty and disjoint intervals of cake. Consider the case that the preferred pieces of cake for two agents are $I_1 = x \cup y$ and $I_2 = y \cup z$. By the non-wasteful property, mechanism M should allocate $x \cup y_1$ to agent 1 and $y_2 \cup z$ to agent 2, where $y_1 \cap y_2 = \emptyset$ and $y_1 \cup y_2 = y$. Without loss of generality, suppose $y_1 \neq \emptyset$. Consider another case that the preferred pieces of cake for two agents are $I_1' = x$ and $I_2' = y \cup z$. Due to non-wasteful property, the mechanism should allocate x to agent 1 and $y \cup z$ to agent 2. Now if agent 1 pretends that his preferred piece is $x \cup y$, he can still obtain $V_1(x \cup y_1) = V_1(x)$ while $V_2(y_2 \cup z)$ is strictly smaller than $V_2(y \cup z)$. Since $\alpha_{1,2} < 0$, agent 1 will get a larger value. The proof is finished by contradiction. □

Note that similar proofs can be constructed when α_{ij} are mixed with positive and negative values. More importantly, the proof does not rely on the formulation of externalities.

Hence, the result applies to a large class of externalities rather than only the model we have. Next, we present the general impossibility result for positive externalities:

Theorem 2. *No truthful, non-wasteful and position independent cake cutting mechanisms exist, even when $\{\alpha_{i,j}\}$ are positive.*

The proof of Theorem 2 relies on the following lemmas.

Lemma 1. *Given two agents with $I_1 = I_2 = X$ where X is a non-empty piece of cake, if $\alpha_{1,2} = \alpha_{2,1} \neq 1$, any non-wasteful and position independent cake cutting mechanism allocates a half of X to each agent.*

Lemma 2. *Given two agents where $I_1 \cap I_2 \neq \emptyset$ and $I_1 \cup I_2 = X$, if for agent 1, $\alpha_{1,2}V_2(I_1 \cap I_2) > V_1(I_1 \cap I_2)$, then any truthful and non-wasteful cake cutting mechanism allocates I_2 to agent 2 and $I_1 \setminus I_2$ to agent 1.*

By Lemma 2, it is easy to observe that if $\alpha_{1,2}V_2(I_1 \cap I_2) > V_1(I_1 \cap I_2)$ and $\alpha_{2,1}V_1(I_1 \cap I_2) > V_2(I_1 \cap I_2)$, then no truthful and non-wasteful cake cutting mechanism exists. But these two conditions induce that $\max(\alpha_{1,2}, \alpha_{2,1}) > 1$ which seems unlikely in real situations. Next we give a stronger impossibility result, which shows that even if we require all $\alpha_{i,j} < 1$, no mechanisms with good properties exist in general.

Proof of Theorem 2 To prove this theorem, we use the instance in Lemma 1. We make $X = I_1 = I_2 = [0, 1]$ and specify $\alpha_{1,2}$ and $\alpha_{2,1}$ such that $1 > \alpha_{1,2} = \alpha_{2,1} > \frac{1}{2}$. It is suggested by Lemma 1 that the entire cake should be split equally to two agents. Now consider the case when agent 1 reports that his valuation density is uniform on $[0, \frac{1}{2} + \epsilon]$ where $\alpha_{2,1} > \frac{1}{2} + \epsilon$. Lemma 2 suggests that the mechanism would allocate $[0, \frac{1}{2} + \epsilon]$ to agent 1 and $[\frac{1}{2} + \epsilon, 1]$ to agent 2. It is clear that agent 1 gets a larger value than before. Hence, this fact suggests that agent 1 could benefit by misreporting $[0, \frac{1}{2} + \epsilon]$ instead of the entire cake $[0, 1]$. Similar arguments apply to multiple agents with $I_1 = \ldots = I_n$ and all $\alpha_{i,j\neq i} = c$ where $\frac{1}{n} < c < 1$. It follows that no truthful, non-wasteful and position independent cake cutting mechanism is possible in general. □

Corollary 1. *Given n agents with all $\frac{1}{n} < \alpha_{i,j\neq i} < 1$, no truthful, non-wasteful and position independent cake cutting mechanisms are possible.*

The results above suggest, when positive externalities exist, in general it is not possible to have truthful cake cutting mechanisms with some good properties. However, it does not rule out the existence of such mechanisms in restricted classes. For example, if all $\alpha_{i,j}$ are equal to 1, meaning the valuation functions of all agents become the same, then a trivial mechanism that allocates the cake to maximize the valuation of all agents solves the problem. The rest of this paper deals with another case in which $\alpha_{i,j\neq i} \leq 1/n^2$.

4 Possibility Result

The intuition behind the proof of Theorem 2 is that agent 1 could benefit by misreporting a smaller preferred piece of cake in the way that agent 2 feels better off if some piece of cake is given to agent 1. As $\alpha_{2,1}$ is high, agent 2 can extract large externalities from agent 1. Although it seems restrictive, it happens in practice. For instance, considering advertisement slots in TV or on the Internet as a cake, a company, which produces and sells parts of vehicles to big automobile companies, would be happier if advertisement slots are assigned to those automobile companies rather than itself. This is because customers watching TV or browsing the Internet have little interest in purchasing individual parts of vehicles. On the other hand, it is often to find that externalities are limited. In this section, we discuss an opposite case when all $\alpha_{i,j}$ are small, i.e., $\alpha_{i,j} \leq \frac{1}{n^2}$. Instead of giving new mechanisms, it is important to understand existing truthful cake cutting algorithms in this new setting. We show a truthful cake cutting algorithm [Chen *et al.*, 2013] performs well even when the externalities exist.

For the sake of completeness, let us present that deterministic cake cutting mechanism as Mechanism 1 and use the following notations. A set of agents is denoted by S, and a piece of cake is denoted by X. Let $D(S, X)$ be the portion of X that is preferred by at least one agent in S, and $avg(S, X) = |D(S, X)|/|S| = |I_S \cap X|/|S|$. An *exact* allocation with respect to S and X allocates each agent in S a piece of cake of length $avg(S, X)$ in his preferred intervals. Given a piece of cake X to be allocated and a set of agents S, Mechanism 1 recursively finds a subset S' of agents with the smallest $avg(S', X)$ and gives an exact allocation $D(S', X)$ to S', then makes a recursive call on the remaining agents $S \setminus S'$ and the remaining cake $X \setminus D(S', X)$. Let $S_1, \ldots, S_m$ be the sequence of agent sets with smallest average chosen by Mechanism 1, and $X_1, \ldots, X_m$ be the sequence of pieces to be allocated in calls to Divide. Note that $avg(S_k, X_k)$ is non-decreasing when k increases. It is shown that Mechanism 1 can be implemented with running time polynomial in n.

Mechanism 1: [Chen *et al.*, 2013]$([0, 1], V_1, \ldots, V_n)$

Divide$(\{1, \ldots, n\}, [0, 1], (V_1, \ldots, V_n))$.
Divide$(S, X, V_1, \ldots, V_n)$
    If $S = \emptyset$, return.
    Let $S_{min} \in \arg\min_{S' \subseteq S} avg(S', X)$, breaking ties arbitrarily;
    Let $E_1, \ldots, E_n$ be an exact allocation with respect to S_{min} and X with arbitrary tie-breaking rule. For each $i \in S_{min}$, set $A_i = E_i$;
    Divide $(S \setminus S_{min}, X \setminus D(S_{min}, X), V_1, \ldots, V_n)$.

Theorem 3. *When there are n agents with all $\alpha_{i,j\neq i} \leq \frac{1}{n^2}$, Mechanism 1 is truthful, non-wasteful, position independent, proportional and swap envy-free.*

Here we would like to emphasize why it is important to understand Mechanism 1 in the setting of externalities. It also gives the motivation that we concentrate on piecewise uniform valuation functions. First, it is very natural to assume that mechanisms that are truthful when externalities exist should also be truthful when no externalities exist. To the best of our knowledge, Mechanism 1 is the only truthful mechanism with other good properties for piecewise uniform valuation functions in the literature. We believe it is definitely worth understanding its performance when externalities exist. Second, in a recent paper [Aziz and Ye, 2014], Aziz and Ye gave two deterministic mechanisms (MEA and CCEA) that are not truthful in general. However, the two mechanisms are truthful for piecewise uniform valuation functions. More importantly, they show that Mechanism 1 is equal to MEA and CCEA in such settings. Therefore, these results make Mechanism 1 a perfect candidate to study when externalities exist.

4.1 Truthfulness

Chen *et al.* [2013] showed that, when Mechanism 1 is used, agents could not get larger pieces in their preferred intervals by misreporting their preferred pieces of cake. This is not enough to guarantee that agents do not have such incentives when externalities are introduced, since agents may strategically report their preferred intervals to benefit by extracting more externalities from the assignments to other agents. It is not trivial to see the truthfulness of Mechanism 1 in this setting. The proof relies on several insights of Mechanism 1, which give us more understanding on it. All lemmas hold without considering the externalities. For this reason, we believe they are also of independent interest.

The first observation of Mechanism 1 establishes the relation between the remaining piece of cake and the already allocated agents.

Observation 1. *When agent i is assigned in some round k during the process of Mechanism 1, i.e. $i \in S_k$, his entire preferred intervals are allocated before or in this round.*

The following two lemmas examine the set of agents allocated in each round. Lemma 3 analyzes the changes on the smallest average length if a subset of agents and their preferred pieces are (recursively) removed. Lemma 4 bounds the ratio on the lengths of preferred intervals between agents in the same S_k (i.e. at the same round). These facts are the key tools to prove the truthfulness of Mechanism 1.

Lemma 3. *Considering a piece X of cake, a set of agents S, for any $S' \subseteq S$ and $avg(S,X) \le avg(S',X)$, it holds $avg(S \setminus S', X \setminus I_{S'}) \le avg(S,X)$.*

Proof.

$$\begin{aligned} avg(S \setminus S', X \setminus I_{S'}) &= \frac{|D(S,X)| - |I_{S'} \cap D(S,X)|}{|S| - |S'|} \\ &= \frac{avg(S,X)|S| - avg(S',X)|S'|}{|S| - |S'|} \\ &\le avg(S,X) \end{aligned}$$

The last inequality is implied by the condition in the lemma. □

Lemma 4. *Let $S_1, \ldots, S_m$ and $X_1, \ldots, X_m$ be defined as above. For any agent $i \in S_k$, if there exists another agent j with $|I_j| < \frac{1}{n}|I_i|$, then $j \in S_{k'}$ where $k' < k$.*

Proof. Let n_k be the number of agents allocated before round k. Because $avg(S_k, X_k)$ is non-decreasing with respect to k, we have

$$\begin{aligned} |I_i \cap X_k| &\ge |I_i| - avg(S_{k-1}, X_{k-1}) \cdot n_k \\ &\ge |I_i| - avg(S_k, X_k) \cdot n_k \end{aligned}$$

Since $i \in S_k$, the smallest average in round k is bounded by assigning $I_i \cap X_k$ to $n - n_k$ agents, that is,

$$avg(S_k, X_k) \ge \frac{|I_i| - avg(S_k, X_k) \cdot n_k}{n - n_k}.$$

This implies $avg(S_k, X_k) \ge \frac{|I_i|}{n}$.

Now let us turn to agent j. In round k, we know that $|I_j \cap X_k| \le |I_j|$. Therefore, setting $S_k = \{j\}$ gives a smaller $avg(S_k, X_k)$, which shows agent j must be allocated before agent i. Then the lemma directly follows. □

By Lemma 3, we show the changes on the assignment for agent j who is previously assigned before or after agent i when agent i misreports his preferred piece of cake.

Lemma 5. *For any agent $i \in N$, let $S_1, \ldots, S_m$ and $X_1, \ldots, X_m$ be defined as above when agent i reports truthfully. Assuming that $i \in S_k$, for any agent $j \in S_{k'}$ with $k' < k$, if agent i misreports his preferred intervals, the value agent j has for his allocated piece of cake is at most the value agent j has for his allocated piece of cake when agent i report truthfully.*

Lemma 6. *For any agent $i \in N$, let $S_1, \ldots, S_m$ and $X_1, \ldots, X_m$ be defined as above when agent i reports truthfully. Assuming that $i \in S_k$, for any agent $j \in S_{k'}$ with $k' \ge k$, if agent i misreports his preferred intervals and the value agent i has for his allocated piece of cake remains the same, then the value agent j has for his allocated piece of cake is at most the value agent j has for his allocated piece of cake when agent i report truthfully.*

Intuitively, the following two lemmas study how the allocations change when the remaining cake changes.

Lemma 7. *Considering a piece $X \subseteq [0,1]$ of cake and a set of agents S, let A be the allocation produced by Mechanism 1. For any for any $S' \subset S$, let A' be the allocation when only the piece $X \setminus I_{S'}$ of cake is assigned to agents in $S \setminus S'$ by Mechanism 1, it holds $|A'_j| \le |A_j|$ for all $j \in S \setminus S'$.*

Lemma 8. *Considering a piece $X \subset [0,1]$ of cake and a set of agents S, let A be the allocation produced by Mechanism 1. For any another piece $X' \subset [0,1]$ of cake with $X \cap X' = \emptyset$, let A' be the allocation when only the piece $X \cup X'$ of cake is assigned to agents in S by Mechanism 1, it holds $|A'_j| \ge |A_j|$ for all $j \in S$.*

Now we are ready to prove the truthfulness of Mechanism 1. The intuition behind the proof is the following. For any agent, we show that he cannot benefit in the way that he extracts more externalities from other agents while keeping a piece of cake with the same value. Furthermore, we show that he cannot benefit from his misreport, in the way that although he gets a less valuable piece of cake than before, he can extract more externalities from other agents which increases his value for the whole allocation eventually.

Property 1. *Mechanism 1 is truthful when all $\alpha_{i,j} \le 1/n^2$.*

Proof. Consider agent i, let A be the allocation produced by Mechanism 1 when he reports truthfully. Recall that the valuation of agent i for allocation A is $V_i(A) = V_i(A_i) + \sum_{j \neq i} \alpha_{i,j} V_j(A_j)$. By the results in [Chen *et al.*, 2013], we know that agent i cannot increase $V_i(A_i)$, that is, the value of agent i for his own piece of cake. It implies that agent i must increase his externalities from other agents by misreporting I'_i in order to benefit. Assume that Mechanism 1 is not truthful, and agent i could benefit from his misreport. Let A' be the allocation produced by Mechanism 1 when agent i misreports, then A' must satisfy one of the following conditions:

1. $V_i(A'_i) = V_i(A_i)$, and there exists another j such that $V_j(A'_j) > V_j(A_j)$;
2. $V_i(A'_i) < V_i(A_i)$ but the increased externalities agent i extracts from the assignments to other agents compensate his loss.

Now we prove that neither of two conditions can happen.

Condition 1**:** Lemma 5 and 6 together show that, when agent i misreports but gets a piece of cake with the same value, any agent cannot get a piece of cake he valued more than that when agent i reports truthfully.

Condition 2**:** To show that agent i fails to extract more externalities from other agents to increase his value for the allocation, we consider the following cases:

- $\forall j \in N \setminus \{i\}, V_j(A'_j) \geq V_j(A_j)$. In this case, only agent i receives a smaller piece of cake, and it implies that every agent in $N \setminus \{i\}$ receives a piece of cake whose length is at most $|A_i| - |A'_i|$ larger than before. Hence, there must exist an agent j such that $|I_j| < \alpha_{i,j}|I_i|$ and $V_j(A'_j) > V_j(A_j)$. It is because the increased externalities agent i extracts from agent j's assignment compensates his loss, which is implied by $\frac{1}{|I_i|} < \frac{\alpha_{i,j}}{|I_j|}$. However, Lemma 4 shows that, if there is an agent j with $|I_j| < \alpha_{i,j}|I_i|$, agent j must be allocated before the round in which agent i is allocated. In addition to this, Lemma 5 shows that, if agent i misreports, agents who are allocated before the round agent i is allocated when agent i reports truthfully, cannot get a larger piece of cake than before. It follows the first case cannot occur.
- Otherwise, besides agent i, there exist other agents who also receive less. It implies that some agent, say j, might get a large piece of cake by an amount more than $|A_i| - |A'_i|$ when agent i misreports. Let us assume $i \in S_k$ when agent i reports truthfully, and $i \in S'_{k'}$ when agent i misreports.

  When $k' > k$, we know that $S'_y = S_y$ for all $y < k$ and agents in $S_1 \cup ... \cup S_{k-1}$ receive the same piece of cake. Agents in $S'_k \cup ... \cup S'_{k'-1}$ receive at most $|I_i \cap A_i| - |I_i \cap A'_i|$ more in total. By Lemma 4, the increased externalities cannot compensate the loss of agent i. Agents after $S'_{k'-1}$ cannot receive more since agent i would compete with them.

  When $k' < k$, we know that that $S_y = S'_y$ for all $y < k'$). It means that agents in $S_1 \cup \ldots \cup S_{k'}$ receive the same pieces of cake as before. By Lemma 5, we know that j does not belong to $S_1 \cup \ldots \cup S_{k-1}$. Therefore, the worst scenario is the following. When agent i misreports, all agents in $S_{k'} \cup \ldots \cup S_{k-1}$ receive weakly less than before and the changes $\sum_{x \in S_{k'} \cup \ldots \cup S_{k-1} \cup \{i\}} |A_x \setminus A'_x|$ are assigned to some agents in $S_k, \ldots, S_m$ excluding agent i. One could verify that it is equivalent to saying that there exists a piece with the length of $\sum_{x \in S_{k'} \cup \ldots \cup S_{k-1} \cup \{i\}} |A_x \setminus A'_x|$ in $I_i \setminus I'_i$ such that this piece is assigned to some agent in $S_k, \ldots, S_m$ excluding agent i. Let $d = avg(S_k, X_k) - avg(S'_{k'}, X'_{k'})$. We know $V_i(A_i) - V_i(A'_i) \geq \frac{d}{|I_i|}$. We first bound the changes on the pieces of cake assigned to agents in $S_{k'} \cup \ldots \cup S_{k-1}$. Since for any $z \in \{k', \ldots, k-1\}$, we know $avg(S_z, X_z) \leq avg(S_k, X_k)$. It implies that the change on the piece of cake assigned to any agent in $S_{k'} \cup \ldots \cup S_{k-1}$ is at most d. Since we have at most $n-1$ agents in $S_{k'} \cup \ldots \cup S_{k-1} \cup \{i\}$, the changes of the pieces of cake assigned to agents in $S_1 \cup \ldots \cup S_{k-1}$ is at most $(n-1)d$. Then let us consider the changes for agents in $(S_k \setminus \{i\}) \cup \ldots \cup S_m$. Let agent g be some agent in $S_1 \cup \ldots \cup S_{k-1}$ who appears last in $S'_1, \ldots, S'_{m'}$. Assume $g \in S'_{t'}$. By Lemma 5, we know that any agent in $(S_k \setminus \{i\}) \cup \ldots \cup S_m \cap (S'_1 \cup \ldots \cup S'_{t'})$ receives weakly less than before. Note that the changes from these agents will not be allocated to agents in $(S_k \setminus \{i\}) \cup \ldots \cup S_m \setminus (S'_1 \cup \ldots \cup S'_{\max\{t',k'\}})$ since all their preferred intervals are allocated before or in the round they are selected. In order to bound the lengths of the extra pieces of cake given to agents in $(S_k \setminus \{i\}) \cup \ldots \cup S_m \setminus (S'_1 \cup \ldots \cup S'_{\max\{t',k'\}})$, we define three scenarios of allocations for those agents. In scenario 1, they have a piece of cake $A_{(S_k \setminus \{i\}) \cup \ldots \cup S_m \setminus (S'_1 \cup \ldots S'_{\max\{t',k'\}})}$, which corresponds to the case when agent i does not misreport; In scenario 2, they have a remaining cake $X'_{\max\{t',k'\}+1} \setminus I_i$, which corresponds to the case when i misreports and also takes away the original I_i; In scenario 3, they have a remaining cake $X'_{\max\{t',k'\}+1}$, which corresponds to the actual case when i misreports. We aim to show that from scenario 1 to 3, every agent in $(S_k \setminus \{i\}) \cup \ldots \cup S_m \setminus (S'_1 \cup \ldots \cup S'_{\max\{t',k'\}})$ receives at most $|I_i \setminus I'_i|$ more. The whole change can be split into two phases. From Scenario 1 to 2, because the available cake becomes strictly less, by Lemma 7, every agent in $(S_k \setminus \{i\}) \cup \ldots \cup S_m \setminus (S'_1 \cup \ldots \cup S'_{\max\{t',k'\}})$ gets weakly less cake. From Scenario 2 to 3, an extra piece of cake with a length of $|I_i - I'_i|$ is added and by Lemma 8, every agent receives weakly more. Combining the two processes we know that every agent receives weakly more but they get at most $|I_i - I'_i|$ more in total. Therefore, the best case for agent i is that a piece of cake with the length of $(n-1)d$ is assigned to an agent v with the biggest $\alpha_{i,v}$ among agents in $(S_k \setminus \{i\}) \cup \ldots \cup S_m \setminus (S'_1 \cup \ldots \cup S'_{\max\{t',k'\}})$. By Lemma 4, we know $|I_v|$ is at least $|I_i|/n$. Hence, the increased externalities agent i could extract from agent

v are at most $\alpha_{i,v}\frac{(n-1)nd}{|I_i|}$. Since $V_i(A) - V_i(A') \geq \frac{d}{|I_i|}$ and $\alpha_{i,v} < \frac{1}{n^2}$, it follows that agent i cannot benefit.

The statement directly follows by the above two cases. □

4.2 Other Properties of Mechanism 1

Chen *et al.* [2013] showed that, without externalities, each agent prefers his own piece of cake (i.e., $V_i(A_i) \geq V_i(A_j)$) and his value of the piece of cake is at least $1/n$. By these facts, when the externalities exist, one could verify that Mechanism 1 is non-wasteful, position independent, proportional and swap envy-free. We also note that, assuming that all $\alpha_{i,j\neq i} \leq 1/n$ and there exists an allocation that optimizes every agent's value, Mechanism 1 outputs such an allocation.

5 Discussion and Conclusion

This paper sheds light on mechanism design in the context of cake cutting when externalities exist. These externalities are benefits or costs to agents from each other's allocation. Specifically, the externalities of agents are modeled as a certain percentage of the value that other agents have for their pieces of cake. Even for this restricted class of externalities, we provide impossibility results. In addition, we show that, when the percentages are small, the truthful mechanism proposed by Chen *et al.* [2013] maintains its truthfulness along with good properties. The results suggest an intuitive interpretation: the existence of deterministic truthful cake cutting mechanisms with good properties relies on the values of $\{\alpha_{i,j}\}$ in the model. The paper leaves an interesting open question — whether there exists a truthful cake cutting mechanism with good properties when some $\alpha_{i,j}$ are greater than $1/n^2$ but not all $\alpha_{i,j}$ are greater than $1/n$. For randomized mechanisms, one could easily see that the randomized mechanism for piecewise linear valuation functions in [Chen *et al.*, 2013] remains truthful in our setting. However, such randomized mechanisms are not non-wasteful. This work also motivates the studies of truthful mechanisms on other models of externalities.

Acknowledgements

We thank the anonymous reviewers for their many insightful comments and suggestions. Minming Li was supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China [Project No. CityU 117913]. Jialin Zhang was supported by National Natural Science Foundation of China (61170062, 61222202, 61433014). Qiang Zhang was supported by FET IP project MULTIPLEX 317532.

References

[Aziz and Ye, 2014] Haris Aziz and Chun Ye. Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *Proceedings of the 10th Conference on Web and Internet Economics (WINE)*, pages 1–14, 2014.

[Balkanski *et al.*, 2014] Eric Balkanski, Simina Brânzei, David Kurokawa, and Ariel D Procaccia. Simultaneous cake cutting. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 566–572, 2014.

[Baumol, 1972] William J Baumol. On taxation and the control of externalities. *The American Economic Review*, 62(3):307–322, 1972.

[Bei *et al.*, 2012] Xiaohui Bei, Ning Chen, Xia Hua, Biaoshuai Tao, and Endong Yang. Optimal proportional cake cutting with connected pieces. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1263–1269, 2012.

[Brânzei and Miltersen, 2013] Simina Brânzei and Peter Bro Miltersen. Equilibrium analysis in cake cutting. In *Proceedings of the 12th International Conference on Autonomous Agents and MultiagentSsystems (AAMAS)*, pages 327–334, 2013.

[Brânzei *et al.*, 2013] Simina Brânzei, Ariel D Procaccia, and Jie Zhang. Externalities in cake cutting. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 55–61, 2013.

[Cachon and Lariviere, 2005] Gérard P Cachon and Martin A Lariviere. Supply chain coordination with revenue-sharing contracts: strengths and limitations. *Management science*, 51(1):30–44, 2005.

[Cachon, 1999] Gérard P Cachon. Competitive supply chain inventory management. In *Quantitative models for supply chain management*, pages 111–146. Springer, 1999.

[Caragiannis *et al.*, 2011] Ioannis Caragiannis, John K Lai, and Ariel D Procaccia. Towards more expressive cake cutting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 127–132, 2011.

[Chen *et al.*, 2013] Yiling Chen, John K Lai, David C Parkes, and Ariel D Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77:284–297, 2013.

[Cohler *et al.*, 2011] Yuga J Cohler, John K Lai, David C Parkes, and Ariel D Procaccia. Optimal envy-free cake cutting. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*, pages 626–631, 2011.

[Haghpanah *et al.*, 2013] Nima Haghpanah, Nicole Immorlica, Vahab Mirrokni, and Kamesh Munagala. Optimal auctions with positive network externalities. *ACM Transactions on Economics and Computation*, 1(2):13, 2013.

[Kurokawa *et al.*, 2013] David Kurokawa, John K Lai, and Ariel D Procaccia. How to cut a cake before the party ends. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, pages 555–561, 2013.

[Netessine and Zhang, 2005] Serguei Netessine and Fuqiang Zhang. Positive vs. negative externalities in inventory management: Implications for supply chain design. *Manufacturing & Service Operations Management*, 7(1):58–73, 2005.

[Nisan and Ronen, 1999] Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC)*, pages 129–140, 1999.

[Procaccia, 2013] Ariel D Procaccia. Cake cutting: not just child's play. *Communications of the ACM*, 56(7):78–87, 2013.

[Robertson and Webb, 1998] Jack Robertson and William Webb. *Cake-cutting algorithms: Be fair if you can*. A K Peters/CRC Press, 1998.

[Segal-Halevi *et al.*, 2015] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. Envy-free cake-cutting in two dimensions. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.