

Truthful Fair Division without Free Disposal

Xiaohui Bei¹, Guangda Huzhang¹, Warut Suksompong²

¹ School of Physical and Mathematical Sciences, Nanyang Technological University

² Department of Computer Science, Stanford University

xhbei@ntu.edu.sg, ghuzhang001@e.ntu.edu.sg, warut@cs.stanford.edu

Abstract

We study the problem of fairly dividing a heterogeneous resource, commonly known as cake cutting and chore division, in the presence of strategic agents. While a number of results in this setting have been established in previous works, they rely crucially on the *free disposal* assumption, meaning that the mechanism is allowed to throw away part of the resource at no cost. In the present work, we remove this assumption and focus on mechanisms that always allocate the entire resource. We exhibit a truthful envy-free mechanism for cake cutting and chore division for two agents with piecewise uniform valuations, and we complement our result by showing that such a mechanism does not exist when certain additional assumptions are made. Moreover, we give truthful mechanisms for multiple agents with restricted classes of valuations.

1 Introduction

Given a heterogeneous divisible resource and a set of interested agents with potentially differing valuations on different parts of the resource, how can we allocate the resource to the agents in such a way that all agents perceive the resulting allocation as fair? The resource is often modeled as a cake in the literature, and the problem, which therefore commonly goes by the name of *cake cutting*, has occupied the minds of mathematicians, computer scientists, economists, and political scientists alike for the past seventy years [Steinhaus, 1948; Brams and Taylor, 1996; Robertson and Webb, 1998; Moulin, 2004; Procaccia, 2016]. Cake in the cake cutting problem is used to represent a *desirable* resource; all agents wish to maximize the amount of resource that they receive. In contrast, the dual problem to cake cutting, known as *chore division*, aims to allocate an *undesirable* resource to the agents, with every agent wanting to receive as little of the resource as possible. Though several algorithms for cake cutting also apply to chore division, the theoretical properties of the two problems differ in many cases, and much less work has been done for chore division than for cake cutting [Peterson and Su, 1998; 2002; Heydrich and van Stee, 2015; Farhadi and Hajiaghayi, 2017; Dehghani *et al.*, 2018].

Perhaps the simplest and most well-known fair division protocol is the *cut-and-choose protocol*, which works for both cake cutting and chore division with two agents. The protocol operates by letting the first agent divide the resource into two parts that she values equally, and letting the second agent choose the part that she prefers. The resulting allocation is always *envy-free*—each agent likes her part at least as much as the other agent’s part, and *proportional*—both agents find their part better than or equal to half of the entire resource. However, the protocol has the disadvantage that it is not *truthful*, meaning that a strategic agent can sometimes benefit from misreporting her valuation to the protocol. For example, if the first agent values the whole cake equally, according to the protocol she will divide the cake into half and get half of her value for the entire cake. However, if she knows that the second agent only cares about the leftmost quarter of the cake, she can divide the cake into the leftmost quarter and the rest, knowing that the second agent will choose the left part and leave her with three-quarters of the cake. The failure to satisfy truthfulness renders the protocol difficult to participate in, since the first agent needs to guess the second agent’s valuation in order to find a beneficial manipulation.

This issue was first addressed by Chen *et al.* [2013], who gave a truthful deterministic cake cutting mechanism that is Pareto optimal, envy-free, and proportional for any number of agents with piecewise uniform valuations. Chen *et al.*’s result shows that fairness and truthfulness are compatible in the allocation of heterogeneous resources. Nevertheless, their result hinges upon a pivotal assumption known as the *free disposal* assumption, which says that the mechanism is allowed to throw away part of the resource without incurring any cost.¹ While certain resources such as cake or machine processing time may be easy to get rid of, for other resources this is not the case. For instance, when we divide a piece of land among antagonistic agents or countries, we cannot simply throw away part of the land, and any piece of land left unallocated constitutes a potential subject of future dispute. The free disposal assumption is even less reasonable when it comes to chore allocation—indeed, with this assumption we might as well simply dispose of the whole chore altogether!

¹Note that free disposal does not preclude Pareto optimality. The mechanism can throw away parts of the resource not valued by any agent and still maintain Pareto optimality; this is exactly what Chen *et al.*’s mechanism does.

With this motivation in mind, we consider in the present paper the problem of fairly and truthfully dividing heterogeneous resources *without the free disposal assumption*. Not having the ability to throw away part of the resource makes the task of the mechanism more complicated. The reason is that even if the mechanism is only allowed to throw away parts that are not valued by any agent, this already prevents agents from gaining by not reporting parts of the resource that no other agent values, in the hope of getting those parts for free along with a larger share of the remaining parts. As Chen *et al.* [2013] noted, getting rid of the free disposal assumption adds “significant complexity” to the problem, since the mechanism would have to specify exactly how to allocate parts that no agent desires. The same group of authors also gave an example illustrating that removing the assumption can be problematic even in the special case of two agents with very simple valuations. Indeed, could it be that there is an impossibility result once we dispose of free disposal?

1.1 Our Results

Throughout the paper, we focus on deterministic mechanisms that are required to allocate the entire resource, which we model as an interval $[0, 1]$. We assume that agents have *piecewise uniform* valuations, meaning that for each agent the cake can be partitioned into desired and undesired intervals, and the agent has the same marginal utility for any fractional piece of any desired interval. We investigate the compatibility of truthfulness and fairness in this setting.

First, in Section 3 we exhibit a truthful, envy-free, and Pareto optimal cake cutting mechanism for two agents (Theorem 1). At a high level, the mechanism lets the two agents “eat” their desired intervals of the cake at the same speed but starting from different ends of the cake. Using a simple reduction from chore division to cake cutting, we also derive a chore division mechanism for two agents with the same set of properties (Theorem 2).

Next, in Section 4 we show that if we add certain requirements for the mechanism on top of being fair and truthful, then no desirable mechanism exists even for two agents. In particular, the impossibility holds when we make any one of the following assumptions in addition to truthfulness and envy-freeness: (i) *anonymity*—the mechanism must treat all agents equally (Theorem 3); (ii) *connected piece assumption*—the mechanism must allocate a single interval to each agent (Theorem 4); and (iii) *position obliviousness*—the values that the agents receive depend only on the lengths of the pieces desired by various subsets of agents and not on the positions of these pieces (Theorem 5).

Finally, in Section 5 we consider the more general setting where there are multiple agents. We assume that each agent only values a single interval of the form $[0, x_i]$. We present a truthful, envy-free, and Pareto optimal cake cutting mechanism (Theorem 6) and a truthful, proportional, and Pareto optimal chore division mechanism (Theorem 7) for any number of agents with valuations in this class.

1.2 Related Work

Cake cutting has been a central topic in the area of fair division and social choice for decades. While the existence

and computation of fair allocations have been extensively studied [Dubins and Spanier, 1961; Stromquist, 1980; Brams and Taylor, 1995; Su, 1999; Aziz and Mackenzie, 2016a; 2016b], the work of Chen *et al.* [2013] that we mentioned earlier was the first to consider incentive issues. Like Chen *et al.*, Maya and Nisan [2012] considered piecewise uniform valuations and gave a characterization of truthful and Pareto optimal mechanisms for two agents. Recently, Alijani *et al.* [2017] presented a truthful envy-free mechanism in the setting where every agent values only a single interval.

For valuation functions beyond piecewise uniform, most results are negative. For example, for piecewise constant valuations, Aziz and Ye [2014] showed that there is no truthful and robust proportional mechanism, two works [Menon and Larson, 2017; Bei *et al.*, 2017] showed that there is no proportional mechanism that allocates connected pieces or is non-wasteful, and Bei *et al.* [2017] also showed that there is no truthful mechanism that satisfies position obliviousness. In the Robertson-Webb query model, Kurokawa *et al.* [2013] showed that there is no truthful envy-free mechanism with bounded queries, while Brânzei and Miltersen [2015] proved that any deterministic truthful mechanism for two agents must be a dictatorship.

In all of the works above, either the free disposal assumption is made, or it is assumed that every piece of the cake is valuable for at least one agent. In contrast, in our work the mechanism is required to always allocate the entire cake. Finally, all aforementioned results are restricted to deterministic mechanisms. If one allows randomization, several truthful in expectation mechanisms that guarantee either proportionality or envy-freeness have been proposed [Mossel and Tamuz, 2010; Chen *et al.*, 2013; Brânzei and Miltersen, 2015].

2 Preliminaries

We consider a heterogeneous divisible resource, which we represent by the interval² $[0, 1]$. A *piece* of the resource is a finite union of disjoint intervals. The resource is to be allocated to n agents a_1, a_2, \dots, a_n . Each agent a_i has a *density function* $f_i : [0, 1] \mapsto \mathbb{R}^+ \cup \{0\}$, which captures how the agent values different parts of the resource. We assume that the agents have *piecewise uniform* valuations, i.e., for each agent a_i , the density function f_i takes on the value 1 on a finite set of intervals and 0 on the remaining intervals. The value of agent a_i for a subset $S \subseteq [0, 1]$ is defined as $v_i(S) = \int_S f_i dx$, which is equivalent to the total length of the intervals in S on which f_i takes on the value 1.³ Let $W_i \subseteq [0, 1]$ denote the piece on which $f_i = 1$. We refer to a setting with agents and their density functions as an *instance*.

An *allocation* of the resource is denoted by a vector $A = (A_1, A_2, \dots, A_n)$, where A_i is a union of finitely many intervals that represents the piece of the resource allocated to a_i , and $A_i \cap A_j = \emptyset$ for any $i \neq j$. We consider two different types of resources: desirable resources, which we represent by a *cake*, and undesirable resources, which we rep-

²Sometimes we will denote the resource by an arbitrary interval $[a, b]$ for simplicity; this can be easily normalized back to $[0, 1]$.

³In some papers, valuations are normalized so that $v_i([0, 1]) = 1$ for all i . We do not follow this convention.

resent by a *chore*. We refer to the problem of allocating the two types of resources as *cake cutting* and *chore division* respectively. The agents want to maximize their value for their piece in cake cutting and minimize this value in chore division. Two fairness properties that we consider are envy-freeness and proportionality. In cake cutting, we say that an allocation (A_1, A_2, \dots, A_n) is *envy-free* if for every agent a_i , we have $v_i(A_i) \geq v_i(A_j)$ for any j . In other words, a_i cannot obtain a larger value from any other agent's share. The allocation is said to be *proportional* if $v_i(A_i) \geq v_i([0, 1])/n$. Envy-freeness and proportionality are defined analogously in chore division but with the inequality signs reversed.

A *mechanism* is a function $\mathcal{M} : (f_1, f_2, \dots, f_n) \mapsto (A_1, A_2, \dots, A_n)$. That is, given the input density functions of the agents, the mechanism computes an allocation for the agents. We only consider deterministic mechanisms in this paper, meaning that the allocation is completely determined by the input density functions. Moreover, we assume that the mechanism has to allocate the entire resource to the agents, i.e., in any allocation (A_1, A_2, \dots, A_n) returned by the mechanism, $\bigcup_{i=1}^n A_i = [0, 1]$. In other words, the mechanism does not have *free disposal*. Note that when the entire resource is allocated, envy-free implies proportionality, and both notions are equivalent in the case of two agents.

We end this section by defining a number of properties of mechanisms that we consider in this paper. Given a vector of input density functions $\mathbf{f} = (f_1, f_2, \dots, f_n)$, let $L_{\mathbf{f}}$ be the indicator function that maps \mathbf{f} to a vector with 2^n components, where each component corresponds to a distinct subset of agents and the value of the component is the length of the piece desired only by that subset of agents.

Definition 1. A mechanism $\mathcal{M} : (f_1, f_2, \dots, f_n) \mapsto (A_1, A_2, \dots, A_n)$ is said to satisfy

- envy-freeness, if it always returns an envy-free allocation;
- proportionality, if it always returns a proportional allocation;
- truthfulness, if it is a dominant strategy for every agent to report her true density function;
- Pareto optimality, if for any allocation returned by the mechanism, there does not exist another allocation that makes no agent worse off and at least one agent better off with respect to the same density functions;
- the connected piece assumption, if each A_i is always a single interval;
- anonymity, if the following holds: For any density functions f_1, f_2, \dots, f_n and any permutation σ of $(1, 2, \dots, n)$, if $\mathcal{M}(f_1, f_2, \dots, f_n) = (A_1, A_2, \dots, A_n)$ and $\mathcal{M}(f_{\sigma(1)}, f_{\sigma(2)}, \dots, f_{\sigma(n)}) = (A'_1, A'_2, \dots, A'_n)$, then $v_i(A_i) = v_i(A'_{\sigma^{-1}(i)})$ for every i .
- position obliviousness, if the following holds: For any vectors of density functions \mathbf{f} and \mathbf{f}' such that $L_{\mathbf{f}} = L_{\mathbf{f}'}$, if $\mathcal{M}(\mathbf{f}) = (A_1, A_2, \dots, A_n)$ and $\mathcal{M}(\mathbf{f}') = (A'_1, A'_2, \dots, A'_n)$, then $v_i(A_i) = v'_i(A'_i)$ for every i .⁴

⁴This is a weaker notion of position obliviousness than the one

Intuitively, a mechanism is anonymous if the utilities that the agents receive do not depend on the identities of the agents, and position oblivious if the values that the agents receive depend only on the lengths of the pieces desired by various subsets of agents and not on the positions of these pieces.

3 Truthful Mechanisms for Two Agents

In this section, we focus on the case of two agents. We show that in this case, there exists a truthful, envy-free, and Pareto optimal mechanism for both cake cutting and chore division, for two agents with arbitrary piecewise uniform valuations.

We first describe the cake cutting mechanism.

Mechanism 1 (for cake cutting between two agents)

Step 1: Find the smallest value of $x \in [0, 1]$ such that $v_1([0, x]) = v_2([x, 1])$.⁵

Step 2: Assign to a_1 the intervals in $[0, x]$ valued by a_1 and the intervals in $[x, 1]$ not valued by a_2 , and assign the rest of the cake to a_2 .

While this is a succinct description of the mechanism, it turns out that the description is somewhat difficult to work with. We next provide an alternative formulation that is more intuitive and will help us in establishing the claimed properties of the mechanism.

Mechanism 1 (alternative formulation)

Phase 1: Let a_1 start at point 0 of the cake moving to the right and a_2 start at point 1 of the cake moving to the left. Let both agents “eat” the cake with the same constant speed, jumping over any interval for which they have no value according to their reported valuations. If the agents are at the same point while both are still eating, go to Phase 3. Else, one of the agents has no more valued interval to eat; go to Phase 2.

Phase 2: Assume that a_i is the agent who has no more valued interval to eat. Let a_i stop and a_{3-i} continue eating. If the agents are at the same point (either while a_{3-i} eats or while a_{3-i} jumps over an interval of zero value), go to Phase 3. Else, both agents have stopped but there is still unallocated cake between their current points. In this case, let a_{3-i} continue eating the unallocated cake until he is at the same point as a_i , and go to Phase 3.

Phase 3: Assume that both agents are at point x of the cake. (It is possible that the two agents meet while both

considered by Bei *et al.* [2017]: Our definition only requires that the agents get the same value if the indicator function of their density functions remain the same, whereas Bei *et al.*'s definition also requires the pieces to be allocated in “equivalent” ways.

⁵The existence of x is guaranteed by the intermediate value theorem.

of them are jumping. In this case, we let a_2 jump first.) Assign any unallocated interval to the left of x to a_2 and any unallocated interval to the right of x to a_1 .

Theorem 1. *Mechanism 1 is a truthful, envy-free, and Pareto optimal cake cutting mechanism for two agents.*

Proof. We begin with truthfulness. Note that there is no incentive for an agent to report an interval that she does not value, since this can only result in the agent wasting time eating such intervals. So the only potential deviation is for the agent to report a strict subset of the intervals that she values. If the agent does not report intervals that she values, then the intervals that she jumps over before the agents meet will be lost to the other agent, and the agent can use the extra time gained from not reporting these intervals to eat intervals of no more than the same length. Moreover, not reporting intervals after the agents meet has no effect on the outcome of the mechanism.

Next, for envy-freeness, it suffices to show that each agent gets at least half of her valued intervals allocated in each phase. In Phase 1, each agent only gains intervals that she values, and loses intervals that she values (due to the other agent's eating) at no more than the same speed. In Phase 2, the agent who continues eating can only gain more, while the agent who has stopped eating has no more interval that she values. In Phase 3, a_1 has no unallocated interval to the left of x that she values, so she cannot lose any unallocated interval that she values. The same argument holds for a_2 .

Finally, our mechanism allocates any interval valued by at least one agent to an agent who values it. This establishes Pareto optimality. \square

Mechanism 1 gives rise to a dual mechanism for two-agent chore division that satisfies the same set of properties.

Mechanism 2 (for chore division between two agents)

Step 1: Use Mechanism 1 to find an initial allocation of the chore, treating the chore valuations as cake valuations.

Step 2: Swap the pieces of the two agents in the allocation from Step 1.

Theorem 2. *Mechanism 2 is a truthful, envy-free, and Pareto optimal chore division mechanism for two agents.*

Proof. First, truthfulness holds because minimizing the chore in the swapped allocation is equivalent to maximizing the chore in the initial allocation, and Theorem 1 shows that this is exactly what Mechanism 1 incentivizes the agents to do. Next, envy-freeness holds again by Theorem 1 because getting at most half of the chore in the swapped allocation is equivalent to getting at least half of the chore in the initial allocation. Finally, in the initial allocation any interval of the chore valued by only one agent is allocated to that agent, so

in the swapped allocation the interval is allocated to the other agent, implying that the mechanism is Pareto optimal. \square

Besides truthfulness, envy-freeness, and Pareto optimality, how do Mechanisms 1 and 2 fare with respect to the other properties defined in Section 2?

- Mechanism 1 is not *anonymous*: If $W_1 = [0, 0.5]$ and $W_2 = [0, 1]$ then both agents get value 0.5, while if $W_1 = [0, 1]$ and $W_2 = [0, 0.5]$ then a_1 gets value 0.75 and a_2 gets value 0.25.
- It is also not *position oblivious*: If $W_1 = [0, 0.5]$ and $W_2 = [0, 1]$ then both agents get value 0.5, while if $W_1 = [0.5, 1]$ and $W_2 = [0, 1]$ then a_1 gets value 0.25 and a_2 gets value 0.75.
- The allocation when $W_1 = [0, 1]$ and $W_2 = [0, 0.5]$ shows that the mechanism does not satisfy the *connected piece assumption*.

The same examples demonstrate that Mechanism 2 likewise satisfies none of the three properties. As we show in the next section, these negative results are in fact not restricted to the two mechanisms that we consider here, but rather apply to all possible cake cutting and chore division mechanisms.

4 Impossibility Results

In this section, we present a number of impossibility results on the existence of fair and truthful mechanisms that satisfy certain additional properties, for both cake cutting and chore division. Interestingly, all of the impossibility results cease to hold if the mechanism is not required to allocate the entire resource, which again highlights the crucial difference that the free disposal assumption makes.

We begin with anonymity. One might expect that a “fair” mechanism should treat the agents equally regardless of their identity. However, the following result shows that anonymity is incompatible with truthfulness and envy-freeness.

Theorem 3. *There does not exist a truthful, envy-free, and anonymous cake cutting mechanism for two agents, even when each agent values a single interval of the form $[0, x_i]$.*

Proof. Suppose that such a mechanism exists. Let $x \in [0, 1]$ and $W_1 = W_2 = [0, x]$. Assume without loss of generality that in this instance, a_1 gets an interval containing point x and ending at point $x + f(x) > x$, possibly among other intervals. By envy-freeness, both agents must get half of the interval $[0, x]$.

If $W_1 = [0, x + \epsilon]$ for some $\epsilon \in [0, f(x)]$ and $W_2 = [0, x]$, then a_1 must get the entire interval $[x, x + \epsilon]$ and half of the interval $[0, x]$. This is because a_2 must get at least half of the interval $[0, x]$, and if a_1 gets less than the whole interval $[x, x + \epsilon]$, she can manipulate by reporting $W_1 = [0, x]$ and getting the whole interval $[x, x + \epsilon]$. By anonymity, if $W_1 = [0, x]$ and $W_2 = [0, x + \epsilon]$ for some $\epsilon \in [0, f(x)]$, a_2 must also get the whole interval $[x, x + \epsilon]$ and half of the interval $[0, x]$.

Now suppose that $W_1 = W_2 = [0, x + \epsilon]$ for some $\epsilon \in [0, f(x)]$. Both agents must get half of the interval $[0, x + \epsilon]$. If a_1 gets more than half of the interval $[x, x + \epsilon]$, then a_2

gets more than half of the interval $[0, x]$. In this case, if $W_2 = [0, x]$, a_2 can manipulate by reporting $W_2 = [0, x + \epsilon]$. So a_1 cannot get more than half of the interval $[x, x + \epsilon]$. By symmetry, neither can a_2 . This means that both agents get exactly half of the interval $[x, x + \epsilon]$. In other words, for any $y \in [x, x + f(x)]$, if $W_1 = W_2 = [0, y]$, then both agents get exactly half of the interval $[x, y]$.

Next, consider the set

$$A := \{(x, y) \in \mathbb{R}_{[0,1]} \times \mathbb{Q}_{[0,1]} \mid x < y < x + f(x)\}.$$

This set is uncountable, since for each of the uncountably many x 's, there is at least one y such that $(x, y) \in A$. If for each y there only exist a finite number of x 's such that $(x, y) \in A$, this set would be countable, which we know is not the case. Hence there exists a y such that $(x, y) \in A$ for infinitely many x 's. Fix such a y .

Finally, suppose that $W_1 = W_2 = [0, y]$. For any of the infinitely many x 's such that $(x, y) \in A$, both agents must receive exactly half of the interval $[x, y]$. However, if the mechanism divides the interval $[0, y]$ into k intervals in the allocation, then there can be at most one value of x per interval, and therefore at most k values in total, with this property. Since k is finite, this gives us the desired contradiction. \square

We remark that *with* the free disposal assumption, Chen et al.'s mechanism is a truthful, envy-free, and anonymous cake cutting mechanism for two agents with arbitrary piecewise uniform valuations. The same authors showed that a particular extension of their mechanism, which allocates the desired pieces of the cake in the same way as their mechanism and allocates the undesired pieces of the cake in a certain simple way, is not truthful [Chen et al., 2013, p. 296]. Since any mechanism that allocates the desired pieces of the cake in this way is also anonymous, Theorem 3 shows that *no* extension of Chen et al.'s mechanism can be truthful.

Next, we turn to the connected piece assumption.

Theorem 4. *There does not exist a truthful and envy-free cake cutting mechanism for two agents that satisfies the connected piece assumption, even when each agent values a single interval of the form $[0, x_i]$.*

Proof. Suppose that such a mechanism exists. First, consider the instance where $W_1 = W_2 = [0, x]$ for some $x \in (0, 1)$. One agent will get the interval $[0, x/2]$ and the other agent the interval $[x/2, 1]$; assume without loss of generality that a_1 gets $[0, x/2]$ and a_2 gets $[x/2, 1]$. Next, consider the instance where $W_1 = [0, x]$ and $W_2 = [0, y]$ for some $y \in (x, 1)$. Then a_2 must still get the interval $[x/2, 1]$; otherwise she can report $W_2 = [0, x]$ instead.

Now, consider the instance where $W_1 = W_2 = [0, y]$. As before, one agent will get the interval $[0, y/2]$ and the other agent the interval $[y/2, 1]$. If a_1 gets $[0, y/2]$, then in the previous instance a_1 can gain more by reporting $W_1 = [0, y]$. Hence it must be that a_2 gets $[0, y/2]$ and a_1 gets $[y/2, 1]$. This means that in the instance where both agents report $[0, y]$, the ordering of the allocated pieces is reversed from the allocation in the instance where both agents report $[0, x]$. Since this holds for any $y > x$, if we take some $z > y$ (obviously $z > x$ also), we find that no allocation works when both agents report $[0, z]$, a contradiction. \square

Bei et al. [2017] showed that a similar impossibility result holds even *with* the free disposal assumption, but using the larger class of piecewise constant valuations. For the class of valuations that we consider in Theorem 4, there exists a simple truthful and envy-free mechanism that always returns a connected allocation assuming free disposal. The mechanism works as follows: Assume that agent a_i declares $W_i = [0, x_i]$ for $i = 1, 2$. If $x_1 \geq x_2$, allocate the interval $[x_1/2, x_1]$ to a_1 and $[0, x_1/2]$ to a_2 ; otherwise allocate the interval $[0, x_2/2]$ to a_1 and $[x_2/2, x_2]$ to a_2 . One can check that this mechanism satisfies the claimed properties.

We now consider position obliviousness and show the nonexistence of a truthful, envy-free, and position oblivious cake cutting mechanism for two agents. In fact, we prove a more general statement that holds for any even number of agents and also uses the weaker notion of proportionality.

Theorem 5. *Let $n = 2k$ for some positive integer k . There does not exist a truthful, proportional, and position oblivious cake cutting mechanism for n agents.*

Proof. Suppose that such a mechanism exists. Assume that the cake is represented by the interval $[0, 4k^2 + k]$.

First, consider the instance where $W_{2i-1} = W_{2i} = [i - 1, i]$ for $i = 1, 2, \dots, k$. Since the interval $[k, 4k^2 + k]$ is of length $4k^2$ and there are $2k$ agents, some agent gets value more than $2k - 1$ from the interval. Assume without loss of generality that a_1 is one such agent, and that a_1 gets the interval $[k, 3k - 1]$. Since the mechanism is proportional, a_1 must get value at least $1/2k$ from the interval $[0, 1]$ as well.

Next, consider the instance where $W_1 = [0, 1] \cup [k, 3k - 1]$, $W_2 = [0, 1]$, and $W_{2i-1} = W_{2i} = [i - 1, i]$ for $i = 2, 3, \dots, k$. Agent a_1 must still get value at least $1/2k$ from the interval $[0, 1]$; otherwise she can report $W_1 = [0, 1]$ instead. This means that a_2 gets a total value of at most $1 - 1/2k$ in this instance.

Finally, consider the instance where $W_1 = W_2 = [0, 1] \cup [k, 3k - 1]$ and $W_{2i-1} = W_{2i} = [i - 1, i]$ for $i = 2, 3, \dots, k$. By proportionality, a_2 must receive value at least 1; let $B_2 \subseteq [0, 1] \cup [k, 3k - 1]$ be a piece of length 1 that a_2 receives. If $W_2 = B_2$ while the other W_i 's remain fixed, then since the mechanism is position oblivious, a_2 must get a total value of at most $1 - 1/2k$. However, in that case a_2 can report $W_2 = [0, 1] \cup [k, 3k - 1]$ and receive value 1. This implies that the mechanism is not truthful and yields the desired contradiction. \square

As with the connected piece assumption, Bei et al. [2017] showed a similar negative result for position obliviousness *with* the free disposal assumption but using the larger class of piecewise constant valuations. For piecewise uniform valuations, Chen et al.'s mechanism is truthful, envy-free, and position oblivious under the free disposal assumption.

We end this section by showing that our impossibility results also carry over to chore division. The idea is the same as the one used in Mechanism 2, except that here we use it to establish negative results.

Corollary 1. *There does not exist a truthful and envy-free chore division mechanism for two agents if one of the follow-*

ing conditions is added: (i) anonymity; (ii) connected piece assumption; (iii) position obliviousness.

Proof. If there were a truthful and envy-free chore division mechanism that satisfies one of the additional properties, we could obtain a cake cutting mechanism with the same properties as follows: First, we use the chore division mechanism to compute an initial allocation of the cake, treating the cake valuations as chore valuations. Then we swap the pieces of the two agents in this allocation. However, the existence of a cake cutting mechanism with these properties would contradict one of Theorems 3, 4, and 5, respectively. \square

5 Extensions to Multiple Agents

In this section, we consider the general setting where we allocate the resource among any number of agents. We assume that each agent a_i only values the interval $[0, x_i]$ for some x_i . Such valuations may appear in a scenario where the agents are dividing machine processing time: agent a_i has a deadline x_i for her jobs, so she would like to maximize the processing time she gets before x_i but has no value for any processing time after x_i . We also remark that the example used to illustrate that removing the free disposal assumption can be problematic consists of *two* agents whose valuations belong to this class [Chen *et al.*, 2013, p. 296]. Hence, designing a fair and truthful algorithm is by no means an easy problem even for this valuation class.

We first describe the cake cutting mechanism.

Mechanism 3 (for cake cutting among n agents)

Step 1: If there is one agent left, the agent gets the entire remaining cake. Else, assume that there are $k \geq 2$ agents and length l of the cake left. Find the maximum $x \in [0, l]$ such that agent i values the entire interval $[(i-1)x, ix]$ for all $i = 1, 2, \dots, k$, and allocate the interval $[(i-1)x, ix]$ to agent i .

Step 2: The agent whose right endpoint of her allocated interval coincides with the right endpoint of her valued piece exits the process. If there are more than one such agent, choose the one with the lowest number.⁶

Step 3: Renumber the remaining agents in the same order starting from 1, and relabel the left endpoint of the remaining cake as point 0. Return to Step 1.

Theorem 6. *Let n be any positive integer. Mechanism 3 is a truthful, envy-free, and Pareto optimal cake cutting mechanism for n agents, if each agent only values a single interval of the form $[0, x_i]$.*

Proof. First, for truthfulness, there are two types of manipulation: moving x_i to the left and to the right. Moving x_i to the left can only cause a_i to quit the process early when she could

⁶There must exist at least one such agent, since otherwise the value of x in Step 1 can be increased.

have gained more by staying on. On the other hand, if moving x_i to the right causes the allocation to change in some round of Step 1, the agent can only get less value from the allocated interval as its right endpoint moves past x_i . Moreover, since she has no more valued intervals to the right, she cannot make up for the loss.

Next, for envy-freeness, if an agent is no longer in the process, she has no more piece of value. During the process, in each round all remaining agents receive an interval of the same length. Since each agent values the entire interval that she receives, she does not envy any other agent.

Finally, our mechanism allocates any interval valued by at least one agent to an agent who values it. This establishes Pareto optimality. \square

Unlike in the case of two agents, there is no simple reduction between cake cutting and chore division in the general case. Nevertheless, our next result shows a truthful and proportional chore division mechanism for any number of agents. We were not able to strengthen the proportionality guarantee to envy-freeness and leave it as an interesting open question for future research.

Mechanism 4 (for chore division among n agents)

Step 1: Let a_1 take the piece $[0, x_1/n] \cup [x_1, 1]$. If some other agent has no value on parts of the interval $[0, x_1/n]$, give those parts to the agent. (If there are several such agents, allocate the parts arbitrarily.)

Step 2: Repeat Step 1 with the next agent up to a_{n-1} and the remaining chore; agent a_i takes the leftmost interval with value x_i/n as well as any piece for which she has no value. (If a_i has value less than x_i/n left, she takes the entire remaining chore.)

Step 3: Agent a_n takes all of the remaining chore.

Theorem 7. *Let n be any positive integer. Mechanism 4 is a truthful, proportional, and Pareto optimal chore division mechanism for n agents, if each agent only values a single interval of the form $[0, x_i]$.*

Proof. We begin with truthfulness. First, any agent who has no value on some piece that the mechanism initially allocates to another agent has no incentive not to take the piece. Apart from this, agent a_n has no control over her allocation, so the mechanism is truthful for her. For any other agent, there are two types of manipulation: moving x_i to the left and to the right. Moving x_i to the right can only increase the value of the piece that a_i has to take. If a_i moves x_i to the left by an amount y , she can save a value of at most y/n but has to take a piece of value y at the end. So a_i does not have a profitable manipulation.

We now consider proportionality. Each agent up to a_{n-1} gets a piece of value at most x_i/n . For a_n , we consider two cases. Let $x = \min(x_1, x_2, \dots, x_{n-1})$. If $x_n \leq x$, then each of the first $n-1$ agents takes at least $1/n$ of the interval $[0, x_n]$, so at most $1/n$ of this interval is left for a_n . Else,

we have $x_n > x$. The intervals $[0, (n - 1)x/n]$ and $[x, 1]$ will not be left to a_n , meaning that a_n receives value at most $x/n < x_n/n$.

Finally, our mechanism allocates any interval for which some agent has no value to one such agent. This establishes Pareto optimality. \square

6 Conclusion and Future Work

In this paper, we study the problem of fairly dividing a heterogeneous resource in the presence of strategic agents and demonstrate the powers and limitations of truthful mechanisms in this setting. An immediate question is whether our mechanisms in Section 3 can be generalized to work for any number of agents with piecewise uniform valuations. While our results in Section 5 provide a partial answer to this question, extending to the general setting seems to require a drastically different idea. Indeed, it could also be that there is an impossibility result once we move beyond the case of two agents.

Another direction is to allow agents to have valuations from a larger class. A natural next step would be to consider the class of *piecewise constant* valuations, in which an agent values each interval uniformly but can have different marginal utilities for different intervals. It is not known whether there exists a deterministic truthful envy-free mechanism even for two agents with piecewise constant valuations, either with or without the free disposal assumption. We believe that this is a theoretically intriguing and practically important question that should be resolved in future work.

References

- [Alijani *et al.*, 2017] Reza Alijani, Majid Farhadi, Mohammad Ghodsi, Masoud Seddighin, and Ahman S. Tajik. Envy-free mechanisms with minimum number of cuts. In *AAAI*, 2017.
- [Aziz and Mackenzie, 2016a] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *FOCS*, 2016.
- [Aziz and Mackenzie, 2016b] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for four agents. In *STOC*, 2016.
- [Aziz and Ye, 2014] Haris Aziz and Chun Ye. Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *WINE*, 2014.
- [Bei *et al.*, 2017] Xiaohui Bei, Ning Chen, Guangda Huzhang, Biaoshuai Tao, and Jiajun Wu. Cake cutting: Envy and truth. In *IJCAI*, 2017.
- [Brams and Taylor, 1995] Steven J. Brams and Alan D. Taylor. An envy-free cake division protocol. *American Mathematical Monthly*, 102(1):9–18, 1995.
- [Brams and Taylor, 1996] Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- [Brânzei and Miltersen, 2015] Simina Brânzei and Peter Bro Miltersen. A dictatorship theorem for cake cutting. In *IJCAI*, 2015.
- [Chen *et al.*, 2013] Yiling Chen, John K. Lai, David Parkes, and Ariel D. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77:284–297, 2013.
- [Dehghani *et al.*, 2018] Sina Dehghani, Alireza Farhadi, MohammadTaghi Hajiaghayi, and Hadi Yami. Envy-free chore division for an arbitrary number of agents. In *SODA*, 2018.
- [Dubins and Spanier, 1961] Lester E. Dubins and Edwin H. Spanier. How to cut a cake fairly. *American Mathematical Monthly*, 68(1):1–17, 1961.
- [Farhadi and Hajiaghayi, 2017] Alireza Farhadi and MohammadTaghi Hajiaghayi. On the complexity of chore division. *arXiv preprint arXiv:1710.00271*, 2017.
- [Heydrich and van Stee, 2015] Sandy Heydrich and Rob van Stee. Dividing connected chores fairly. *Theoretical Computer Science*, 593:51–61, 2015.
- [Kurokawa *et al.*, 2013] David Kurokawa, John K. Lai, and Ariel D. Procaccia. How to cut a cake before the party ends. In *AAAI*, 2013.
- [Maya and Nisan, 2012] Avishay Maya and Noam Nisan. Incentive compatible two player cake cutting. In *WINE*, 2012.
- [Menon and Larson, 2017] Vijay Menon and Kate Larson. Deterministic, strategyproof, and fair cake cutting. In *IJCAI*, 2017.
- [Mossel and Tamuz, 2010] Elchanan Mossel and Omer Tamuz. Truthful fair division. In *SAGT*, 2010.
- [Moulin, 2004] Hervé Moulin. *Fair Division and Collective Welfare*. MIT press, 2004.
- [Peterson and Su, 1998] Elisha Peterson and Francis Edward Su. Exact procedures for envy-free chore division. *preprint*, 1998.
- [Peterson and Su, 2002] Elisha Peterson and Francis Edward Su. Four-person envy-free chore division. *Mathematics Magazine*, 75(2):117–122, 2002.
- [Procaccia, 2016] Ariel D. Procaccia. Cake cutting algorithms. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13, pages 311–329. Cambridge University Press, Cambridge, 2016.
- [Robertson and Webb, 1998] Jack Robertson and William Webb. *Cake-Cutting Algorithms: Be Fair if You Can*. Peters/CRC Press, 1998.
- [Steinhaus, 1948] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.
- [Stromquist, 1980] Walter Stromquist. How to cut a cake fairly. *American Mathematical Monthly*, 87(8):640–644, 1980.
- [Su, 1999] Francis Edward Su. Rental harmony: Sperner’s lemma in fair division. *American Mathematical Monthly*, 106(10):930–942, 1999.