## ORIGINAL PAPER

# Truthful fair division without free disposal 

Xiaohui Bei ${ }^{1} \cdot$ Guangda Huzhang $^{1} \cdot$ Warut Suksompong $^{2}$

Received: 16 August 2019 / Accepted: 15 April 2020 / Published online: 25 April 2020 © The Author(s) 2020


#### 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 and 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 constraints are imposed on the mechanisms. Moreover, we provide bounds on the efficiency of mechanisms satisfying various properties, and 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 (Brams and Taylor 1996; Moulin 2004; Procaccia 2016; Robertson and Webb 1998; Steinhaus 1948). 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

[^0]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 on chore division than on cake cutting (Dehghani et al. 2018; Farhadi and Hajiaghayi 2018; Heydrich and van Stee 2015; Peterson and Su 1998, 2002).

Perhaps the simplest and most well-known fair division protocol is the cut-andchoose 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 entire chore altogether!

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. For example, with free disposal allowed, the mechanism can throw away parts that are not valued by any agent, thereby preventing an agent from gaining by not reporting parts of the resource for which she is the only one who has positive value, in the hope of getting some of those parts for free along with a larger share of the remaining parts. As Chen et al. (2013) noted,

[^1]Table 1 Existence of mechanisms satisfying truthfulness in combination with other properties in the case of two agents

| Properties of mechanism | Existence |
| :--- | :--- |
| Truthful + envy-free + Pareto optimal | Yes (Theorems 1 and 2) |
| Truthful + connected + position oblivious | Yes (trivial mechanism) |
| Truthful + anonymous | No (Theorem 3 and Corollary 1) |
| Truthful + connected + positive share | No (Theorem 4 and Corollary 1) |
| Truthful + position oblivious + positive share | No (Theorem 5 and Corollary 1) |

By "trivial mechanism" we refer to a mechanism that allocates the entire resource to one fixed agent. All results hold for both cake cutting and chore division
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 the 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 Sect. 3, we show that truthfulness and fairness are compatible when there are two agents by exhibiting a truthful, envy-free, and Pareto optimal cake cutting mechanism (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).

In Sect. 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 both 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). In fact, our first impossibility result still holds even if we remove envy-freeness, while for our second and third results we can replace envyfreeness by the significantly weaker fairness notion of positive share, which requires
that each agent receive a positive utility (for cake cutting) or not incur the entire cost (for chore division). Our results in Sects. 3 and 4 are summarized in Table 1.

In Sect. 5, we investigate the efficiency of mechanisms satisfying combinations of truthfulness, envy-freeness, and Pareto optimality. Our metric is the efficiency ratio, a standard measure for the welfare performance of mechanisms. ${ }^{2}$ For cake cutting, we show that our mechanism in Sect. 3 achieves an efficiency ratio of 3/4, and that this is optimal among all truthful and envy-free mechanisms (Theorem 7 and Corollary 2). By contrast, for chore division we show that our mechanism achieves the worst possible efficiency ratio of 0 ; however, this is necessarily the case for any truthful and envy-free mechanism (Theorem 10). Moreover, we provide tight bounds on the efficiency ratio of mechanisms satisfying each of envy-freeness and Pareto optimality, for both cake cutting and chore division (Theorems 8, 9, and 11).

Finally, in Sect. 6, we consider the more general setting where there are multiple agents. We assume that each agent only values a single interval of the form $\left[0, x_{i}\right]$. We present a truthful, envy-free, and Pareto optimal cake cutting mechanism (Theorem 12) and a truthful, proportional, and Pareto optimal chore division mechanism (Theorem 13) 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 social choice and economics for decades. While the existence and computation of fair allocations have been extensively studied (Aziz and Mackenzie 2016a, b; Brams and Taylor 1995; Dubins and Spanier 1961; Goldberg et al. 2020; Stromquist 1980; Su 1999), the work of Chen et al. (2013) that we mentioned earlier was the first to consider incentive issues. As with 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 and envy-free mechanism in the setting where every agent only values 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, ${ }^{3}$ two works (Bei et al. 2017; Menon and Larson 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 and envy-free

[^2]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. ${ }^{4}$ Furthermore, 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 (Brânzei and Miltersen 2015; Chen et al. 2013; Mossel and Tamuz 2010).

Finally, preferences analogous to piecewise uniform valuations have been considered in other settings including matching (Bade 2015; Bogomolnaia and Moulin 2004) and collective choice (Aziz et al. 2019; Bogomolnaia et al. 2005; Duddy 2015), where they are known as dichotomous preferences. Fairness notions corresponding to proportionality and positive share have also been studied in the collective choice setting.

## 2 Preliminaries

We consider a heterogeneous divisible resource, which we represent by the interval $[0,1]$. A piece of the resource is a finite union of disjoint intervals. The resource is to be allocated to $n$ agents $a_{1}, a_{2}, \ldots, a_{n}$. Each agent $a_{i}$ has a density function $f_{i}:[0,1] \rightarrow \mathbb{R}^{+} \cup\{0\}$, 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 $v_{i}(S)=\int_{S} f_{i} d x$, which is equivalent to the total length of the intervals in $S$ on which $f_{i}$ takes on the value $1 .{ }^{5}$ For chore division we will refer to the value as cost. 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 $A=\left(A_{1}, A_{2}, \ldots, A_{n}\right)$, where $A_{i}$ is a union of finitely many intervals ${ }^{6}$ that represents the piece of the resource allocated to $a_{i}$, and $A_{i} \cap A_{j}$ has measure zero 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 represent 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 allocated piece in cake cutting and minimize their cost for their allocated piece in chore division.

[^3]We are now ready to define the fairness properties that we consider in this paper.

Definition 1 In cake cutting, we say that an allocation $\left(A_{1}, A_{2}, \ldots, A_{n}\right)$ satisfies

- envy-freeness, if for every agent $a_{i}$, we have $v_{i}\left(A_{i}\right) \geq v_{i}\left(A_{j}\right)$ for any $j$;
- proportionality, if for every agent $a_{i}$, we have $v_{i}\left(A_{i}\right) \geq v_{i}([0,1]) / n$;
- positive share, if for every agent $a_{i}$, we have $v_{i}\left(A_{i}\right)>0$.

In chore division, envy-freeness and proportionality are defined analogously but with the inequality signs reversed, while positive share is defined by the inequality $v_{i}\left(A_{i}\right)<v_{i}([0,1])$.

A mechanism is a function $\mathcal{M}:\left(f_{1}, f_{2}, \ldots, f_{n}\right) \rightarrow\left(A_{1}, A_{2}, \ldots, A_{n}\right)$ which, given the input density functions of the agents, computes an allocation for them. 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 $\left(A_{1}, A_{2}, \ldots, A_{n}\right)$ 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. Moreover, both envy-freeness and proportionality are stronger than positive share.

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 $\mathbf{f}=\left(f_{1}, f_{2}, \ldots, f_{n}\right)$, 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 (and not by any agent outside the subset).

Definition 2 A mechanism $\mathcal{M}:\left(f_{1}, f_{2}, \ldots, f_{n}\right) \rightarrow\left(A_{1}, A_{2}, \ldots, A_{n}\right)$ is said to satisfy

- envy-freeness, if it always returns an envy-free allocation;
- proportionality, if it always returns a proportional allocation;
- positive share, if it always returns an allocation satisfying positive share;
- 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: Let $f_{1}, f_{2}, \ldots, f_{n}$ be any density functions, and $\sigma$ be any permutation of $(1,2, \ldots, n)$. If $\mathcal{M}\left(f_{1}, f_{2}, \ldots, f_{n}\right)=\left(A_{1}, A_{2}, \ldots, A_{n}\right)$ and $\mathcal{M}\left(f_{\sigma(1)}, f_{\sigma(2)}, \ldots, f_{\sigma(n)}\right)=\left(A_{1}^{\prime}, A_{2}^{\prime}, \ldots, A_{n}^{\prime}\right)$, then $v_{i}\left(A_{i}\right)=v_{i}\left(A_{\sigma^{-1}(i)}^{\prime}\right)$ for every $i$.
- position obliviousness, if the following holds: Let $\mathbf{f}$ and $\mathbf{f}^{\prime}$ be any vectors of density functions such that $L_{\mathbf{f}}=L_{\mathbf{f}^{\prime}}$. If $\mathcal{M}(\mathbf{f})=\left(A_{1}, A_{2}, \ldots, A_{n}\right)$ and $\mathcal{M}\left(\mathbf{f}^{\prime}\right)=\left(A_{1}^{\prime}, A_{2}^{\prime}, \ldots, A_{n}^{\prime}\right)$, then $v_{i}\left(A_{i}\right)=v_{i}^{\prime}\left(A_{i}^{\prime}\right)$ for every $i{ }^{7}$

Intuitively, a mechanism is anonymous if the utility 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 $v_{1}([0, x])=v_{2}([x, 1]) .{ }^{8}$
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,

[^4]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 she 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 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 $a_{1}$.

From the perspective of this formulation, Mechanism 1 bears a resemblance to the probabilistic serial mechanism for the random assignment problem (Bogomolnaia and Moulin 2001); both mechanisms let the agents simultaneously eat their desired part of the resource at the same speed.

We now prove the claimed properties of Mechanism 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.

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 that incurs a cost to 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.

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

- Mechanism 1 is not anonymous: If $W_{1}=[0,0.5]$ and $W_{2}=[0,1]$, both agents get value 0.5 , while if $W_{1}=[0,1]$ and $W_{2}=[0,0.5], a_{1}$ gets value 0.75 and $a_{2}$ gets value 0.25 .
- It is also not position oblivious: If $W_{1}=[0,0.5]$ and $W_{2}=[0,1]$, both agents get value 0.5 , while if $W_{1}=[0.5,1]$ and $W_{2}=[0,1], a_{1}$ 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, even when envy-freeness is removed or significantly weakened.

## 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. In fact, one of the results holds even without envy-freeness, while for the remaining results we can replace envy-freeness by the much weaker notion of positive share. 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 also observe that since these results are negative, they also hold for more general classes of valuations beyond piecewise uniform.

Before we introduce our results, some comments on truthful mechanisms are in order. At first glance, truthfulness may appear as a strong requirement that limits
the feasible mechanisms to Mechanisms 1 and 2 and those returning a fixed allocation. However, the examples that we exhibit next show that the class of truthful mechanisms is surprisingly rich. We only describe cake cutting mechanisms, but it is possible to obtain corresponding chore division mechanisms via the reduction used in Mechanism 2. The latter two classes of mechanisms are adaptations of the picking and exchange mechanisms (Amanatidis et al. 2017) to the divisible goods setting.

- A variant of Mechanism 1 where the two agents eat the cake at possibly different speeds fixed in advance. We can also partition the cake into a number of parts, "flip" some of these parts horizontally, "paste" some of the parts together, and run the variant of Mechanism 1 with different speeds for different parts.
- A dictatorship mechanism, where one fixed agent takes an arbitrary superset of her valued piece, and the other agent takes the rest of the cake.
- A picking mechanism, where the cake is cut in advance into a number of parts, some designated to $a_{1}$ and the rest to $a_{2}$. Each agent has a set of offers, each offer corresponding to a subset of the parts designated to her, and the agent picks an offer that she likes most. (For example, if parts 1, 3, 5, and 6 are designated to $a_{1}$, a possible set of offers is $\{\{1,3\},\{3,5\},\{1,5,6\}\}$. In this case, if $a_{1}$ picks the offer $\{3,5\}$, then $a_{2}$ gets parts 1 and 6 .)
- An exchange mechanism, where the cake is cut in advance into a number of parts, some initially allocated to $a_{1}$ and the rest to $a_{2}$. Some exchange deals are considered; each exchange deal involves a subset of $a_{1}$ 's parts and a subset of $a_{2}$ 's parts, where each part appears in at most one exchange deal. An exchange deal materializes if the exchange benefits both agents.

It is also possible to combine these mechanisms by partitioning the cake in advance and running different mechanisms on different parts. Among these classes of mechanisms, only the first class, with the agents eating at the same speed, is envy-free.

We now proceed to our impossibility results, which demonstrate that even though there are several ways to obtain truthfulness, none of them is compatible with fairness and other desirable properties. We begin by showing that anonymity is directly at variance with truthfulness.

Theorem 3 There does not exist a truthful and anonymous cake cutting mechanism for two agents, even when each agent values a single interval of the form $\left[0, x_{i}\right]$.

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 $x$ (perhaps as the left endpoint) and ending at point $x+f(x)>x$, possibly among other intervals. By anonymity, 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 value at least $x / 2+\epsilon$, since otherwise she can manipulate by reporting $W_{1}=[0, x]$. Therefore $a_{2}$ gets value at most $x / 2$ in this instance. By anonymity, if $W_{1}=[0, x]$ and
$W_{2}=[0, x+\epsilon]$ for some $\epsilon \in(0, f(x))$, then $a_{2}$ gets value at least $x / 2+\epsilon$ and $a_{1}$ at most $x / 2$.

Now suppose that $W_{1}=W_{2}=[0, x+\epsilon]$ for some $\epsilon \in(0, f(x))$. By anonymity, 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 [ $\left.0, x\right]$. 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 receive exactly half of the interval $[x, y]$.

Next, consider the set

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

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 some $y$ 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.

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 and show that it is incompatible with truthfulness and positive share.

Theorem 4 There does not exist a truthful cake cutting mechanism for two agents that satisfies positive share and the connected piece assumption, even when each agent values a single interval of the form $\left[0, x_{i}\right]$.

Proof Suppose that such a mechanism exists. First, we fix $x \in(0,1)$ and assume that $W_{1}=[0, x]$. For each $y \in(0,1)$, if $W_{2}=[0, y]$, then by the connected piece assumption, $a_{2}$ gets either a "left interval" of the form $[0, z]$ or a "right interval" of the form [ $z, 1$ ]; by positive share, we have $z \in(0,1)$, so $a_{2}$ never gets an interval that is simultaneously a left interval and a right interval. We claim that there exist $f(x), g(x)>0$ with $f(x)+g(x)<x$ such that if $W_{2}=[0, y]$ for some $0<y<f(x)$, then $a_{2}$ gets a
superset of $[0, y]$, while if $W_{2}=[0, y]$ for some $f(x) \leq y<f(x)+g(x)$, then $a_{2}$ gets exactly $[0, f(x)]$. To prove this claim, we consider two cases.

- Case 1 For some $y \in(0,1]$, if $W_{2}=[0, y]$, then $a_{2}$ gets a right interval. We show that this right interval must be the same for all such $y$. Suppose for contradiction that $a_{2}$ gets $\left[z_{1}, 1\right]$ when $W_{2}=\left[0, y_{1}\right]$ and gets $\left[z_{2}, 1\right]$ when $W_{2}=\left[0, y_{2}\right]$, for some $z_{1}<z_{2}$. By positive share, we have $z_{2}<y_{2}$. When $W_{2}=\left[0, y_{2}\right], a_{2}$ can manipulate by reporting $W_{2}=\left[0, y_{1}\right]$ and obtain a higher utility. So we must have $z_{1}=z_{2}$. Hence there exists $h(x) \in(0,1)$ such that whenever $a_{2}$ gets a right interval, she gets exactly the interval $[h(x), 1]$. In particular, if $W_{2}=[0, y]$ for some $y \in(0, h(x))$, then positive share implies that $a_{2}$ must get a left interval.

Next, define
$f(x):=\sup _{y \in(0,1]}\left\{z \mid a_{2}\right.$ gets the interval $[0, z]$ when $\left.W_{2}=[0, y]\right\}$.
From the previous paragraph, we have $f(x)>0$; moreover, positive share implies that $f(x) \leq x$. If $W_{2}=[0, y]$ for some $y \in(0, f(x)]$, then $a_{2}$ must get the entire interval $[0, y]$ (and possibly more); otherwise, by definition of $f(x)$, she can manipulate to get the more of the interval $[0, y]$. In particular, when $W_{2}=[0, f(x)]$, the definition of $f(x)$ implies that $a_{2}$ gets exactly the interval [ $0, f(x)$ ]. By positive share, this also means that $f(x)<x$.

Now, let $g(x) \in(0, h(x))$ be such that $f(x)+g(x)<x$. Suppose that $W_{2}=[0, y]$ for some $f(x)<y<f(x)+g(x)$. If $a_{2}$ gets a right interval, this interval must be $[h(x), 1]$, which yields value strictly less than $f(x)$ to her. However, she can manipulate by reporting $W_{2}=[0, f(x)]$ and get value $f(x)$. So $a_{2}$ must get a left interval of length at least $f(x)$. By definition of $f(x)$, this must be exactly the interval $[0, f(x)]$. Therefore $f(x)$ and $g(x)$ satisfy the conditions of our claim.

- Case 2 For any $y \in(0,1]$, if $W_{2}=[0, y]$, then $a_{2}$ gets a left interval. We define $f(x)$ as in the previous case and choose any $g(x)>0$ such that $f(x)+g(x)<x$. A similar argument as before shows that $f(x)$ and $g(x)$ satisfy the conditions of our claim.

Next, let $x^{\prime} \in(f(x), x)$. Since our choice of $x$ was arbitrary, for $x^{\prime}$ there also exist $f\left(x^{\prime}\right)$ and $g\left(x^{\prime}\right)$ satisfying analogous conditions. Suppose first that $f\left(x^{\prime}\right)<f(x)$. Let $y \in\left(f\left(x^{\prime}\right), f\left(x^{\prime}\right)+g\left(x^{\prime}\right)\right)$ be such that $y<f(x)$, and assume that $W_{2}=[0, y]$. If $W_{1}=[0, x]$, then since $y<f(x)$, our claim implies that $a_{2}$ gets at least the interval $[0, y]$. However, $a_{1}$ can manipulate by reporting $W_{1}=\left[0, x^{\prime}\right]$ so that $a_{2}$ only gets the interval $\left[0, f\left(x^{\prime}\right)\right]$. So we cannot have $f\left(x^{\prime}\right)<f(x)$. Likewise, if $f\left(x^{\prime}\right)>f(x)$, we can take $y \in(f(x), f(x)+g(x))$ to arrive at a contradiction. Hence we must have $f\left(x^{\prime}\right)=f(x)$. Since our choice of $x^{\prime}$ was arbitrary, this holds for any $x^{\prime} \in(f(x), x)$.

Finally, let $x^{\prime \prime} \in(0, f(x))$. Assume that when $W_{1}=\left[0, x^{\prime \prime}\right]$ and $W_{2}=[0, f(x)]$, $a_{1}$ gets value $r>0$. Let $r^{\prime} \in(0, \min \{r, x-f(x)\})$. If $W_{1}=\left[0, f(x)+r^{\prime}\right]$ and $W_{2}=[0, f(x)]$, then since $f\left(f(x)+r^{\prime}\right)=f(x), a_{2}$ gets exactly the interval [0, f(x)], which leaves value at most $r^{\prime}$ to $a_{1}$. However, in this case $a_{1}$ can manipulate by reporting $W_{1}=\left[0, x^{\prime \prime}\right]$ to obtain value $r>r^{\prime}$. This is a contradiction.

Bei et al. (2017) showed that a similar impossibility result with envy-freeness instead of positive share holds 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 $W_{i}=\left[0, x_{i}\right]$ for $i=1,2$. If $x_{1} \geq x_{2}$, allocate the interval $\left[x_{1} / 2, x_{1}\right]$ to $a_{1}$ and $\left[0, x_{1} / 2\right]$ to $a_{2}$; otherwise allocate the interval $\left[0, x_{2} / 2\right]$ to $a_{1}$ and $\left[x_{2} / 2, x_{2}\right]$ to $a_{2}$. One can check that this mechanism satisfies the claimed properties.

We now consider position obliviousness and show the non-existence of a truthful and position oblivious cake cutting mechanism that satisfies positive share for two agents. In Appendix A, we prove a statement that uses proportionality instead of positive share but holds for any even number of agents.

Theorem 5 There does not exist a truthful and position oblivious cake cutting mechanism for two agents that satisfies positive share.

Proof Suppose that such a mechanism exists. First, we claim that for any $x \in(0,1 / 3)$, if $W_{1}=W_{2}=[0, x]$, then one of the agents gets length strictly less than $x$ from the interval $[x, 1]$. Assume that this is not the case, and suppose that $a_{1}$ receives value $r>0$ in the instance where $W_{1}=W_{2}=[0, x]$. Since both agents get length at least $x$ from $[x, 1]$, we may assume without loss of generality that $a_{1}$ gets $[x, 2 x]$ (perhaps among other pieces). This means that if $W_{1}=[0,2 x]$ and $W_{2}=[0, x], a_{1}$ must receive value at least $x+r$; otherwise she can manipulate by reporting $W_{1}=[0, x]$. We now show that $a_{1}$ must still receive value at least $x+r$ when $W_{1}=W_{2}=[0,2 x]$. This suffices to establish our claim since by symmetry, $a_{2}$ also receives value more than $x$ when $W_{1}=W_{2}=[0,2 x]$, which is impossible.

If it were not the case that $a_{1}$ receives value at least $x+r$ when $W_{1}=W_{2}=[0,2 x]$, then $a_{2}$ receives value more than $x-r$ in this instance. In the instance where $W_{1}=[0,2 x]$, and $W_{2}$ has length $x$ and contains a piece of length more than $x-r$ that $a_{2}$ values in the previous instance, $a_{2}$ must receive value more than $x-r$; otherwise she can manipulate by reporting $W_{2}=[0,2 x]$. Hence $a_{1}$ receives value less than $x+r$ in this instance. However, we know that $a_{1}$ receives value at least $x+r$ when $W_{1}=[0,2 x]$ and $W_{2}=[0, x]$. This is a contradiction with position obliviousness. So $a_{1}$ indeed receives value at least $x+r$ when $W_{1}=W_{2}=[0,2 x]$, proving our claim.

Next, fix $x \in(0,1 / 3)$. Assume without loss of generality that when $W_{1}=W_{2}=[0, x], a_{2}$ gets length strictly less than $x$ from the interval $[x, 1]$, and that $a_{2}$ does not get any piece outside the interval [1-x,1]. Consider the instance where $W_{1}=[0,1-x]$ and $W_{2}=[0, x]$. Note that $a_{1}$ must get value more than $1-2 x$; otherwise she can manipulate by reporting $W_{1}=[0, x]$. Now, consider the instance where $W_{1}=W_{2}=[0,1-x]$. If $a_{2}$ gets value at least $x$, let $B_{2}$ be a piece of length $x$ that $a_{2}$ values and gets. If $W_{1}=[0,1-x]$ and $W_{2}=B_{2}, a_{2}$ must get at least the entire $B_{2}$, and hence $a_{1}$ gets value at most $1-2 x$. However, we know that $a_{1}$ receives value greater than $1-2 x$ when $W_{1}=[0,1-x]$ and $W_{2}=[0, x]$, which contradicts position obliviousness. Hence $a_{2}$ gets value less than $x$ when $W_{1}=W_{2}=[0,1-x]$. This
means that when $W_{1}=[0,1-x]$ and $W_{2}=[0,1], a_{2}$ gets value less than $2 x$; otherwise $a_{2}$ can manipulate by reporting $W_{2}=[0,1]$ when in fact $W_{2}=[0,1-x]$.

Finally, assume that $a_{2}$ receives value $s>0$ when $W_{1}=W_{2}=[0,1]$. Let $s^{\prime}:=s / 3 \in(0,1 / 3)$. If $a_{1}$ manipulates by reporting $W_{1}=\left[0,1-s^{\prime}\right], a_{2}$ gets value less than $2 s^{\prime}<s$. Hence this is a beneficial manipulation for $a_{1}$, a contradiction.

As with the connected piece assumption, Bei et al. (2017) showed a similar negative result for position obliviousness and envy-freeness 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 chore division mechanism for two agents that satisfies any one of the following set of conditions: (i) anonymity; (ii) connected piece assumption and positive share; (iii) position obliviousness and positive share.

Proof If there were a truthful mechanism that satisfies one of the additional set of 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.

## 5 Bounds on the efficiency ratio

In this section, we investigate the efficiency of mechanisms satisfying various properties using the efficiency ratio, which is a standard measure for the welfare performance of mechanisms. In particular, we show that for both cake cutting and chore division, no truthful and envy-free mechanism attains a better efficiency ratio than the mechanisms we presented in Sect. 3.

In contrast to the rest of the paper, in this section we do engage in interpersonal comparison of utility. For this reason, we assume that each agent has a value (or cost) of 1 for the entire resource, and scale the agent's utility for a subset of the resource proportionally. ${ }^{9}$ In cake cutting, the social welfare of an allocation is the sum of the values that the two agents receive, and the efficiency ratio of a mechanism is the infimum, over all feasible instances, of the ratio between the social welfare of the allocation produced by the mechanism and the maximum social welfare.

[^5]Similarly, in chore division, the social cost of an allocation is the sum of the costs that the two agents incur, and the efficiency ratio of a mechanism is the infimum, over all instances, of the ratio between the minimum social cost and the social cost of the allocation produced by the mechanism. ${ }^{10}$ Note that by definition, the efficiency ratio is always at most 1 .

### 5.1 Cake cutting

We start with cake cutting. Our first result gives a lower bound on the efficiency ratio of envy-free and Pareto optimal mechanisms.

Theorem 6 Any envy-free and Pareto optimal cake cutting mechanism for two agents has efficiency ratio at least 3/4, and this bound is tight.

Proof Fix an envy-free and Pareto optimal cake cutting mechanism. Consider an instance where the two agents value pieces of length $x \leq y$ respectively, and these pieces have an overlap of length $z \leq x$. The welfare-maximizing allocation gives the entire overlap to $a_{1}$, yielding welfare $1+\frac{y-z}{y}=2-\frac{z}{y}$. Next, we derive a lower bound for the welfare of the allocation that the mechanism produces. We consider two cases.

- Case 1: $z \leq x / 2$. Since the mechanism is Pareto optimal, the welfare achieved is at least the welfare when the entire overlap is allocated to $a_{2}$. This yields welfare $1+\frac{x-z}{x}=2-\frac{z}{x} \geq \frac{3}{2}$. Hence the efficiency ratio is at least $\frac{3 / 2}{2}=\frac{3}{4}$.
- Case 2: $z>x / 2$. The minimum welfare is achieved by minimizing $a_{1}$ 's share of the overlap. Since the mechanism is envy-free, $a_{1}$ gets a piece of length at least $x / 2$. This means that she gets length at least $\frac{x}{2}-(x-z)=z-\frac{x}{2}$ from the overlap. If she gets length exactly $z-\frac{x}{2}$ from the overlap, the social welfare is

$$
\frac{1}{2}+\frac{y-z+x / 2}{y} \geq \frac{3}{2}-\frac{z}{2 y},
$$

where the inequality holds since $x \geq z$. Hence the efficiency ratio is at least

$$
\frac{\frac{3}{2}-\frac{z}{2 y}}{2-\frac{z}{y}}=\frac{1}{2} \cdot \frac{3-\frac{z}{y}}{2-\frac{z}{y}} \geq \frac{3}{4}
$$

The tightness of the bound follows from the fact that Mechanism 1 has efficiency ratio 3/4, which we show in Corollary 2.

Next, we provide an upper bound on the efficiency ratio of truthful and envy-free mechanisms.

[^6]Theorem 7 Any truthful and envy-free cake cutting mechanism for two agents has efficiency ratio at most $3 / 4$, and this bound is tight.

Proof Fix a truthful and envy-free cake cutting mechanism. Consider the instance where $W_{1}=W_{2}=[0, \epsilon]$ for some small $\epsilon>0$. Since the mechanism is envy-free, it must allocate exactly half of the interval $[0, \epsilon]$ to each agent. Assume without loss of generality that the mechanism gives $a_{1}$ at least half of the interval $[\epsilon, 1]$; let $B_{1} \subseteq[\epsilon, 1]$ be the piece that $a_{1}$ receives. Next, consider the instance where $W_{1}=[0, \epsilon] \cup B_{1}$ and $W_{2}=[0, \epsilon]$. By truthfulness, $a_{1}$ must receive $B_{1}$ and half of the interval $[0, \epsilon]$, and therefore $a_{2}$ can only receive half of $[0, \epsilon]$. For small $\epsilon$, the optimal social welfare approaches 2 , while the social welfare of the allocation produced by the mechanism is at most $1 / 2+1=3 / 2$. Hence the efficiency ratio of the mechanism is at most $3 / 4$, as claimed.

The tightness of the bound follows from the fact that Mechanism 1 has efficiency ratio 3/4, which we show in Corollary 2.

The bounds in Theorems 6 and 7 yield the following corollary, which immediately implies the tightness of both bounds. Moreover, it follows from the corollary that no truthful and envy-free cake cutting mechanism has a better efficiency ratio than Mechanism 1.

Corollary 2 Any truthful, envy-free, and Pareto optimal cake cutting mechanism for two agents has efficiency ratio exactly 3/4. In particular, the efficiency ratio of Mechanism 1 is 3/4.

Next, we prove tight upper and lower bounds for mechanisms satisfying each of envy-freeness and Pareto optimality.

Theorem 8 Any envy-free cake cutting mechanism for two agents has efficiency ratio at least $1 / 2$ and at most $\frac{2+\sqrt{3}}{4} \approx 0.933$, and both bounds are tight.

Proof We consider the lower and upper bounds in turn.
Lower bound An envy-free mechanism always gives each agent a utility of at least $1 / 2$, so the social welfare of the resulting allocation is at least 1 . On the other hand, the optimal social welfare is at most 2 . Hence the efficiency ratio is at least $1 / 2$. To show that this bound is tight, consider a mechanism that always outputs the allocation output by the Mechanism 1, except when $W_{1}=[0,1]$ and $W_{2}=[0, \epsilon]$ for some small $\epsilon>0$. For this instance, the mechanism gives the piece $[\epsilon / 2,(1+\epsilon) / 2]$ to $a_{2}$ and the rest of the cake to $a_{1}$. This is an envy-free allocation, and both agents receive utility $1 / 2$. On the other hand, for small $\epsilon$, the optimal welfare in this instance approaches 2.

Upper bound Caragiannis et al. (2011, Theorem 9) showed that for two agents with arbitrary (not necessarily piecewise uniform) valuations, the maximum welfare of an envy-free allocation is no less than $\frac{2+\sqrt{3}}{4}$ times the optimal social welfare. Hence a mechanism that always returns an envy-free allocation with maximum
welfare has efficiency ratio at least $\frac{2+\sqrt{3}}{4}$. To show that this bound is tight, we adapt the example of Caragiannis et al. to piecewise uniform valuations. Consider the instance where $W_{1}=[0, \sqrt{3}-1]$ and $W_{2}=[0,1]$. The welfare-maximizing allocation gives the entire overlap to $a_{1}$, yielding welfare $1+(2-\sqrt{3})=3-\sqrt{3}$. On the other hand, any envy-free allocation gives $a_{1}$ a piece of length at most $1 / 2$ from the overlap. So the maximum welfare of an envy-free allocation is $\frac{1 / 2}{\sqrt{3}-1}+\frac{1}{2}=\frac{3+\sqrt{3}}{4}$. Hence the efficiency ratio of any envy-free mechanism is at most $\frac{(3+\sqrt{3}) / 4}{3-\sqrt{3}}=\frac{2+\sqrt{3}}{4}$, as claimed.

Theorem 9 Any Pareto optimal cake cutting mechanism for two agents has efficiency ratio at least $1 / 2$ and at most 1 , and both bounds are tight.

Proof We consider the lower and upper bounds in turn.
Lower bound Consider an instance where the two agents value pieces of length $x \leq y$ respectively, and these pieces have an overlap of length $z \leq x$. The welfaremaximizing allocation gives the entire overlap to $a_{1}$, yielding welfare $1+\frac{y-z}{y}=2-\frac{z}{y}$. On the other hand, the lowest welfare of a Pareto optimal allocation is achieved by giving the entire overlap to $a_{2}$, which yields welfare $1+\frac{x-z}{x}=2-\frac{z}{x}$. Hence the efficiency ratio of a Pareto optimal mechanism is at least $\frac{2-z / x}{2-z / y} \xrightarrow{x} \frac{2-1}{2-0}=\frac{x}{2}$. To show that this bound is tight, consider a mechanism that always outputs the allocation output by Mechanism 1, except when $W_{1}=[0,1]$ and $W_{2}=[0, \epsilon]$ for some small $\epsilon>0$. For this instance, the mechanism gives the entire cake to $a_{1}$. This is a Pareto optimal allocation with social welfare 1 . On the other hand, for small $\epsilon$, the optimal welfare approaches 2.

Upper bound By definition, the efficiency ratio of any mechanism cannot exceed 1. Moreover, a mechanism that always returns an allocation with maximum social welfare is Pareto optimal and has efficiency ratio exactly 1.

### 5.2 Chore division

We now address chore division and exhibit a surprising contrast to cake cutting. Our first result shows that unlike for cake cutting, for chore division no truthful and envy-free mechanism has a positive efficiency ratio. Since Mechanism 2 is truthful and envy-free, the result applies to the mechanism as well.

Theorem 10 Any truthful and envy-free chore division mechanism for two agents has efficiency ratio 0 . In particular, the efficiency ratio of Mechanism 2 is 0 .

Proof Fix a truthful and envy-free chore division mechanism. Consider the instance where $W_{1}=W_{2}=[0, \epsilon]$ for some small $\epsilon>0$. Since the mechanism is envy-free, it must allocate exactly half of the interval $[0, \epsilon]$ to each agent. Assume without loss of generality that the mechanism gives $a_{2}$ at least half of the interval $[\epsilon, 1]$;
let $B_{2} \subseteq[\epsilon, 1]$ be the piece that $a_{2}$ receives. Next, consider the instance where $W_{1}=[0, \epsilon] \cup B_{2}$ and $W_{2}=[0, \epsilon]$. By truthfulness, $a_{1}$ must receive at most half of the interval $[0, \epsilon]$, and therefore $a_{2}$ receives at least half of the interval $[0, \epsilon]$. For small $\epsilon$, the minimum social cost in this instance approaches 0 , while the social cost of the allocation produced by the mechanism is at least $1 / 2$. Hence the efficiency ratio of the mechanism is 0 , as claimed.

Theorem 10 also implies that any combination of truthfulness, envy-freeness, and Pareto optimality does not suffice to guarantee a positive efficiency ratio. On the other hand, a mechanism that always returns an allocation with minimum social cost is Pareto optimal and has the highest possible efficiency ratio of 1 . Our next result shows that an envy-free mechanism can achieve an efficiency ratio of up to $8 / 9$.

Theorem 11 Any envy-free chore division mechanism for two agents has efficiency ratio at most $8 / 9$, and this bound is tight.

Proof Caragiannis et al. (2011, Theorem 17) showed that for two agents with arbitrary (not necessarily piecewise uniform) valuations, the minimum cost of an envyfree allocation is no more than $9 / 8$ times the optimal social cost. Hence a mechanism that always returns an envy-free allocation with minimum cost has efficiency ratio $8 / 9$. To show that this bound is tight, we adapt the example of Caragiannis et al. to piecewise uniform valuations. Consider the instance where $W_{1}=[0,2 / 3]$ and $W_{2}=[0,1]$. The cost-minimizing allocation gives the entire overlap to $a_{2}$, yielding cost $2 / 3$. On the other hand, any envy-free allocation gives $a_{2}$ a piece of length at most $1 / 2$ from the overlap. So the minimum cost of an envy-free allocation is $\frac{1 / 6}{2 / 3}+\frac{1}{2}=\frac{3}{4}$. Hence the efficiency ratio of any envy-free mechanism is at most $\frac{2 / 3}{3 / 4}=\frac{8}{9}$, as claimed.

## 6 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 (or has cost on) the interval $\left[0, x_{i}\right]$ 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 $\ell$ of the cake left. Find the maximum $x \in[0, \ell]$ such that agent $i$ values the entire interval $[(i-1) x, i x]$ for all $i=1,2, \ldots, k$, and allocate the interval $[(i-1) x, i x]$ 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. ${ }^{11}$
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 12 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 $\left[0, x_{i}\right]$.

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 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.

We remark that constructing truthful and envy-free mechanisms that work beyond the class of valuations in Theorem 12 appears to be highly nontrivial. For instance, even if every agent only values a single interval (not necessarily starting at 0 ), then a mechanism that tries to find valued intervals of equal length according to the agent ordering no longer works: if $W_{1}=[0.5,1]$ and $W_{2}=[0,0.5]$, such a mechanism would not be able to allocate any of the cake. Alternatively, one could try a generalization of Mechanism 3 that finds the smallest $x$ such that $i x$ is the right endpoint of some agent $i$ 's valued interval. However, if $W_{1}=W_{2}=[0.5,1]$, then this mechanism would allocate $[0,0.5]$ to agent 1 and $[0.5,1]$ to agent 2 , thereby leaving agent 1 envious.

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. ${ }^{12}$ We

[^7]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 $\left[0, x_{1} / n\right] \cup\left[x_{1}, 1\right]$. If some other agent has no cost on parts of the interval $\left[0, x_{1} / n\right]$, give those parts to the agent. (If there are several such agents, allocate the parts arbitrarily.)
Step 2 Proceed similarly with the next agent up to $a_{n-1}$ and the remaining chore; agent $a_{i}$ takes the leftmost interval with value $y_{i} / n$ as well as any piece for which she has no cost, where $y_{i}:=\min \left(x_{1}, x_{2}, \ldots, x_{i}\right)$. (If $a_{i}$ has cost less than $y_{i} / n$ left, she takes the entire remaining chore.)
Step 3 Agent $a_{n}$ takes all of the remaining chore.
Theorem 13 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 has cost on a single interval of the form $\left[0, x_{i}\right]$.

Proof We begin with truthfulness. First, any agent who has no cost 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 while staying to the right of $y_{i-1}=\min \left(x_{1}, x_{2}, \ldots, x_{i-1}\right)$, nothing changes. Else, she moves $x_{i}$ by an amount $z$ on the left of $y_{i-1}$. In this case, she can save a cost of at most $z / n$ but has to take a piece of cost $z$ at the end. So $a_{i}$ does not have a profitable manipulation.

We now consider proportionality. Each agent $a_{i}$ up to $a_{n-1}$ gets a piece of cost at most $y_{i} / n \leq x_{i} / n$. For $a_{n}$, we consider two cases. If $x_{n} \leq y_{n-1}$, then each of the first $n-1$ agents takes at least $1 / n$ of the interval $\left[0, x_{n}\right]$, so at most $1 / n$ of this interval is left for $a_{n}$. Else, we have $x_{n}>y_{n-1}$. The intervals $\left[0,(n-1) y_{n-1} / n\right]$ and $\left[y_{n-1}, 1\right]$ will not be left to $a_{n}$, meaning that $a_{n}$ incurs a cost of at most $y_{n-1} / n<x_{n} / n$.

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

## 7 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 that arises is whether the mechanisms in Sect. 3 can be generalized to work for any number of agents with
piecewise uniform valuations. While our results in Sect. 6 provide a partial answer to this question, extending to the general setting seems to require a drastically different idea. Indeed, while other examples of truthful mechanisms given in Sect. 4 can be generalized to multiple agents, these extensions do not satisfy envy-freeness or even positive share. Of course, it could also be that there is an impossibility result once we move beyond the case of two agents.

Another direction for future work 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. Intriguingly, it is not known whether there exists a deterministic truthful and envy-free mechanism even for two agents with piecewise constant valuations, either with or without the free disposal assumption. Again, the question is still open even if we relax envy-freeness to positive share; it does not seem clear whether (significantly) relaxing the fairness notion helps in designing a truthful mechanism for piecewise constant valuations.

Finally, our results in Sect. 5 leave open the question of determining the best efficiency ratio that can be achieved by a truthful mechanism. This quantity has been called the price of truthfulness by Maya and Nisan (2012), who studied it in the setting where free disposal is allowed. While we have established tight upper bounds on the efficiency ratio of truthful and envy-free mechanisms for both cake cutting and chore division, our current techniques are not sufficient to show that the bounds remain tight in the absence of envy-freeness. If a truthful mechanism with a higher efficiency ratio were to exist, it would be useful in situations where only resistance to strategic behavior is desired.

Acknowledgements This work was partially supported by the European Research Council (ERC) under grant number 639945 (ACCORD) and by a Stanford Graduate Fellowship. A preliminary version of the paper appeared in Proceedings of the 27th International Joint Conference on Artificial Intelligence. The authors thank Felix Brandt, Dominik Peters, and Christian Stricker for helpful discussions, and the anonymous reviewers for valuable comments.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licen ses/by/4.0/.

## A Additional impossibility result

Theorem 14 Let $n=2 k$ 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. For simplicity, we represent the cake by the interval $\left[0,4 k^{2}+k\right]$; this can be easily normalized back to $[0,1]$.

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

Next, consider the instance where $W_{1}=[0,1] \cup[k, 3 k-1], W_{2}=[0,1]$, and $W_{2 i-1}=W_{2 i}=[i-1, i]$ for $i=2,3, \ldots, k$. Agent $a_{1}$ must still get value at least $1 / 2 k$ 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 $1-1 / 2 k$ in this instance.

Finally, consider the instance where $W_{1}=W_{2}=[0,1] \cup[k, 3 k-1]$ and $W_{2 i-1}=W_{2 i}=[i-1, i]$ for $i=2,3, \ldots, k$. By proportionality, $a_{2}$ must receive value at least 1 ; let $B_{2} \subseteq[0,1] \cup[k, 3 k-1]$ 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 $1-1 / 2 k$. However, in that case $a_{2}$ can report $W_{2}=[0,1] \cup[k, 3 k-1]$ and receive value 1 . This implies that the mechanism is not truthful and yields the desired contradiction.

## References

Alijani R, Farhadi M, Ghodsi M, Seddighin M, Tajik AS (2017) Envy-free mechanisms with minimum number of cuts. In: Proceedings of the 31st AAAI conference on artificial intelligence, pp 312-318
Amanatidis G, Birmpas G, Markakis E (2016) On truthful mechanisms for maximin share allocations. In: Proceedings of the 25th international joint conference on artificial intelligence, pp 31-37
Amanatidis G, Birmpas G, Christodoulou G, Markakis E (2017) Truthful allocation mechanisms without payments: characterization and implications on fairness. In: Proceedings of the 18th ACM conference on economics and computation, pp 545-562
Aziz H, Mackenzie S (2016a) A discrete and bounded envy-free cake cutting protocol for any number of agents. In: Proceedings of the 57th annual symposium on foundations of computer science, pp 416-427
Aziz H, Mackenzie S (2016b) A discrete and bounded envy-free cake cutting protocol for four agents. In: Proceedings of the 48th annual ACM SIGACT symposium on theory of computing, pp 454-464
Aziz H, Ye C (2014) Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In: Proceedings of the 10th international conference on web and internet economics, pp 1-14
Aziz H, Bogomolnaia A, Moulin H (2019) Fair mixing: the case of dichotomous preferences. In: Proceedings of the 20th ACM conference on economics and computation, pp 753-781
Bade S (2015) Multilateral matching under dichotomous preferences. Working paper
Bei X, Chen N, Huzhang G, Tao B, Wu J (2017) Cake cutting: envy and truth. In: Proceedings of the 26th international joint conference on artificial intelligence, pp 3625-3631
Bei X, Lu X, Manurangsi P, Suksompong W (2019) The price of fairness for indivisible goods. In: Proceedings of the 28th international joint conference on artificial intelligence, pp 81-87
Bogomolnaia A, Moulin H (2001) A new solution to the random assignment problem. J Econ Theory 100(2):295-328
Bogomolnaia A, Moulin H (2004) Random matching under dichotomous preferences. Econometrica 72(1):257-279
Bogomolnaia A, Moulin H, Stong R (2005) Collective choice under dichotomous preferences. J Econ Theory 122(2):165-184

Brams SJ, Taylor AD (1995) An envy-free cake division protocol. Am Math Mon 102(1):9-18
Brams SJ, Taylor AD (1996) Fair division: from cake-cutting to dispute resolution. Cambridge University Press, Cambridge
Brânzei S, Miltersen PB (2015) A dictatorship theorem for cake cutting. In: Proceedings of the 24th international joint conference on artificial intelligence, pp 482-488
Caragiannis I, Kaklamanis C, Kanellopoulos P, Kyropoulou M (2011) The efficiency of fair division. Theory Comput Syst 50(4):589-610
Chen Y, Lai JK, Parkes DC, Procaccia AD (2013) Truth, justice, and cake cutting. Games Econ Behav 77:284-297
Dehghani S, Farhadi A, Hajiaghayi M, Yami H (2018) Envy-free chore division for an arbitrary number of agents. In: Proceedings of the 29th annual ACM-SIAM symposium on discrete algorithms, pp 2564-2583
Dubins LE, Spanier EH (1961) How to cut a cake fairly. Am Math Mon 68(1):1-17
Duddy C (2015) Fair sharing under dichotomous preferences. Math Soc Sci 73:1-5
Farhadi A, Hajiaghayi M (2018) On the complexity of chore division. In: Proceedings of the 27th international joint conference on artificial intelligence, pp 226-232
Goldberg PW, Hollender A, Suksompong W (2020) Contiguous cake cutting: Hardness results and approximation algorithms. In: Proceedings of the 34th AAAI conference on artificial intelligence (forthcoming)
Heydrich S, van Stee R (2015) Dividing connected chores fairly. Theor Comput Sci 593:51-61
Kurokawa D, Lai JK, Procaccia AD (2013) How to cut a cake before the party ends. In: Proceedings of the 27th AAAI conference on artificial intelligence, pp 555-561
Maya A, Nisan N (2012) Incentive compatible two player cake cutting. In: Proceedings of the 8th international workshop on internet and network economics, pp 170-183
Menon V, Larson K (2017) Deterministic, strategy proof, and fair cake cutting. In: Proceedings of the 26th international joint conference on artificial intelligence, pp 352-358
Mossel E, Tamuz O (2010) Truthful fair division. In: Proceedings of the 3rd international symposium on algorithmic game theory, pp 288-299
Moulin H (2004) Fair division and collective welfare. MIT Press, Cambridge
Peterson E, Su FE (1998) Exact procedures for envy-free chore division. Working paper
Peterson E, Su FE (2002) Four-person envy-free chore division. Math Mag 75(2):117-122
Procaccia AD (2016) Cake cutting algorithms. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) Handbook of computational social choice, chapter 13. Cambridge University Press, Cambridge, pp 311-329
Robertson J, Webb W (1998) Cake-cutting algorithms: be fair if you can. Peters/CRC Press, Boca Raton
Steinhaus H (1948) The problem of fair division. Econometrica 16(1):101-104
Stromquist W (1980) How to cut a cake fairly. Am Math Mon 87(8):640-644
Su FE (1999) Rental harmony: Sperner's lemma in fair division. Am Math Mon 106(10):930-942
Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.


[^0]:    Warut Suksompong
    warut.suksompong@cs.ox.ac.uk
    1 Nanyang Technological University, Singapore, Singapore
    2 University of Oxford, Oxford, UK

[^1]:    ${ }^{1}$ Note 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.

[^2]:    ${ }^{2}$ The efficiency ratio is closely related to the price of fairness, which measures the worst-case welfare loss due to imposing fairness constraints (Bei et al. 2019; Caragiannis et al. 2011; Heydrich and van Stee 2015).
    ${ }^{3}$ Aziz and Ye (2014) call an allocation robust proportional if it satisfies the following property: even if an agent perturbs her value density function, as long as the ordinal information of the function is unchanged, then the allocation remains proportional. Robust proportionality is a stronger notion than standard proportionality.

[^3]:    ${ }^{4}$ On the other hand, in the study of truthful mechanisms for allocating indivisible items, it is commonly assumed that the mechanism is required to allocate all items (Amanatidis et al. 2017, 2016).
    ${ }^{5}$ In some papers, valuations are normalized so that $v_{i}([0,1])=1$ for all $i$. We do not follow this convention except in Sect. 5.
    ${ }^{6}$ We assume without loss of generality that these intervals are closed intervals. The assumption that each $A_{i}$ is a union of finitely many intervals is standard in the cake-cutting literature (see, e.g., Procaccia 2016).

[^4]:    ${ }^{7}$ This is a weaker notion of position obliviousness than the one 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.
    ${ }^{8}$ The existence of $x$ is guaranteed by the intermediate value theorem.

[^5]:    ${ }^{9}$ For this to make sense, we also assume that each agent has a positive value (or cost) on a nonempty subset of the interval $[0,1]$.

[^6]:    ${ }^{10}$ Here we assume that $\frac{0}{0}=1$. This is a reasonable choice because if the mechanism produces an allocation with zero social cost for a certain instance, there is no loss of efficiency in that instance.

[^7]:    ${ }^{11}$ There always exists at least one such agent, since otherwise the value of $x$ in Step 1 can be increased.
    ${ }^{12} \mathrm{We}$ are grateful to Bo Li for pointing out an error in an earlier version of this mechanism.

