

Deterministic, Strategyproof, and Fair Cake Cutting

Vijay Menon and Kate Larson

David R. Cheriton School of Computer Science

University of Waterloo

{vijay.menon, kate.larson}@uwaterloo.ca

Abstract

We study the classic cake cutting problem from a mechanism design perspective, in particular focusing on deterministic mechanisms that are strategyproof and fair. We begin by looking at mechanisms that are non-wasteful and primarily show that for even the restricted class of piecewise constant valuations there exists no direct-revelation mechanism that is strategyproof and even approximately proportional. Subsequently, we remove the non-wasteful constraint and show another impossibility result stating that there is no strategyproof and approximately proportional direct-revelation mechanism that outputs contiguous allocations, again, for even the restricted class of piecewise constant valuations. In addition to the above results, we also present some negative results when considering an approximate notion of strategyproofness, show a connection between direct-revelation mechanisms and mechanisms in the Robertson-Webb model when agents have piecewise constant valuations, and finally also present a (minor) modification to the well-known Even-Paz algorithm that has better incentive-compatible properties for the cases when there are two or three agents.

1 Introduction

Imagine a scenario where there is heterogeneous divisible good that is to be divided among a certain set of n agents. For an appropriately chosen notion of fairness, ideally, we would like this resource to be divided fairly among these n agents and so a natural question that arises is on how one would accomplish this. The cake-cutting problem metaphorically refers to this very scenario and it represents a fundamental problem in the theory of fair division. More formally, in the cake cutting problem, the cake is modelled as the interval $[0, 1]$ and each of the n agents is assumed to have a valuation function over the cake. The goal, as described above, is to partition the cake so that it is fair according to some chosen notion of fairness. Starting with the work of Steinhaus [1948], it has been studied over the last several decades by mathematicians, economists, and political scientists (see e.g. the books by Brams and Taylor [1996] and Robertson

and Webb [1998]), and more recently has attracted the attention of computer scientists (see e.g. the survey by Procaccia [2016]) partly due to the nature of the challenges involved—that often require an algorithm design or complexity point of view—and partly due to its relevance in the design of multi-agent systems [Chevalleyre *et al.*, 2006].

Most work in the cake cutting literature focuses on addressing the above described scenario and so most of it is concerned with computing a fair allocation using proportionality or envy-freeness as the main fairness criterion (see [Aziz and Mackenzie, 2016], and also [Procaccia, 2016] for a recent survey). While addressing it is indeed the core part of the problem, the fact that the valuation functions of the agents may be their private information entails that the designed protocols or mechanisms be strategyproof—meaning that there is no incentive to misreport one’s valuation function—for if otherwise the agents can lie about their valuation functions and potentially benefit from it. Therefore, the focus of this paper, like in [Chen *et al.*, 2010; 2013; Mossel and Tamuz, 2010; Maya and Nisan, 2012; Aziz and Ye, 2014; Brânzei and Miltersen, 2015], is to look at mechanisms that are not only fair, but also strategyproof.

In this paper, we look at deterministic mechanisms and we are primarily focused on the direct-revelation model—one where the agents reveal their entire valuation function. Our main objective is to better understand the limits of determinism when it comes to imposing strategyproofness and fairness (or approximate notions of either of them)—i.e., to see what is or is not achievable with deterministic mechanisms—and to this end we make the following contributions.

- a) In Section 3 we begin by looking at non-wasteful mechanisms and we show a strong impossibility which in turn strengthens an impossibility result given [Aziz and Ye, 2014, Theorem 7]. In particular, we prove (in Theorem 1) that for any $n \geq 2$ and $0 \leq \epsilon_1, \epsilon_2 < \frac{1}{n}$ such that $3\epsilon_1 + \epsilon_2 < \frac{1}{n}$, there is no deterministic and non-wasteful mechanism for n agents with piecewise constant valuations that is ϵ_1 -strategyproof and ϵ_2 -proportional.
- b) In Section 4 we remove the non-wasteful constraint and provide three main results. The first one (Theorem 5) shows that for any $n \geq 2$ and $0 \leq \epsilon < \frac{1}{n}$, there is no deterministic mechanism for n agents with piecewise constant valuations that makes contiguous alloca-

tions and is strategyproof and ϵ -proportional. The second result (Proposition 6) establishes a connection between direct-revelation mechanisms and mechanisms in the Robertson-Webb model for the case when agents have piecewise constant valuations. And finally, the third one (Theorem 8) is a positive result where we present a (minor) modification to the well-known Even-Paz algorithm that has better incentive-compatible properties for the cases when there are two or three agents.

1.1 Related Work

The cake cutting problem has been studied using two input models—Robertson-Webb model and the direct-revelation model. Therefore, we can classify the work on strategic aspects of cake cutting into two classes based on the input model used and look at each of them separately.

Among related papers that operate in the direct-revelation model [Chen *et al.*, 2010; 2013; Mossel and Tamuz, 2010; Maya and Nisan, 2012; Aziz and Ye, 2014; Li *et al.*, 2015; Alijani *et al.*, 2017], the ones which are most relevant to results in this paper are the works of Chen *et al.* [2010; 2013] and Aziz and Ye [2014]. Chen *et al.* [2010; 2013] were the first to look at strategyproof cake cutting and their main result was a deterministic, strategyproof, envy-free, proportional, and Pareto-optimal mechanism for the case when the agents have piecewise uniform valuations. Additionally, they also presented randomized algorithms that are truthful in expectation, proportional, and envy free for piecewise linear valuations.

The other most relevant work that uses the same model is that of Aziz and Ye [2014]. They present two deterministic mechanisms, CCEA and MEA, for piecewise constant valuations, where CCEA is robust envy-free¹ and non-wasteful, and MEA is Pareto-optimal and envy-free. Although both of them generalize the mechanism proposed by Chen *et al.* [2010; 2013], they do not remain strategyproof for piecewise constant valuations. Additionally, they also showed that there is no mechanism that is strategyproof, robust proportional¹, and non-wasteful for piecewise constant valuations.

Among papers that use the Robertson-Webb model, the two that are most related are the works of Kurokawa *et al.* [2013] and Brânzei and Miltersen [2015]. Kurokawa *et al.* [2013] studied envy-free cake cutting and their main result was a parameterized protocol for piecewise linear valuations. Additionally, they also showed that there is no mechanism of complexity bounded only by a function of the number of agents and the total number of breakpoints that is strategyproof and envy-free for piecewise constant valuations.

The other most relevant paper that operates in the Robertson-Webb model is that of Brânzei and Miltersen [2015]. Brânzei and Miltersen [2015] showed that any deterministic and strategyproof protocol is a dictatorship when there are only two hungry agents and for the case when there are more than two hungry agents they showed that there is at least one agent who gets an empty piece. Additionally, they

¹Robust envy-freeness and robust proportionality are much stronger notions than envy-freeness and proportionality, respectively; see [Aziz and Ye, 2014] for more details.

also provided a randomized mechanism that is strategyproof in expectation and ϵ -proportional for any $\epsilon > 0$.

Finally, we also briefly remark on the predefined divisible goods setting which has been extensively studied in fair division (see, for instance, [Cheung, 2016] and the references therein). In the predefined goods setting, the task is to fairly allocate m divisible items among n agents, where the private information of an agent is the value she has for each of the items. So, in effect, the predefined goods setting can be considered as a special case of the setting of cake cutting with piecewise constant valuations—i.e., one where the valuation functions of all the agents have the same set of breakpoints—and therefore all the negative results in that setting carry over to the cake cutting setting. However, some of the nice positive results with regards to strategyproof mechanisms in this setting (see, e.g., [Cole *et al.*, 2013] and [Cheung, 2016]) do not carry over to the cake cutting setting as here, informally, the difficulty in achieving strategyproofness arises from the fact that a mechanism has to guard against two types of manipulations: i) with regards to the breakpoints in the valuation functions and ii) with regards to the values assigned to the pieces between two consecutive breakpoints.

2 Preliminaries

The cake—which is a heterogeneous divisible good—is modelled as the unit interval $[0, 1]$. A piece of cake X is a finite union of disjoint (except at the boundaries of the intervals) subintervals of $[0, 1]$. An interval is denoted by I and the length of an interval $I = [x, y]$ is given by $|I| = y - x$. In any cake cutting instance, there are n agents who want a share of the cake and we denote them by $[n] = \{1, \dots, n\}$. Each agent $i \in [n]$ has a non-negative, integrable, and private value density function (also referred to as their utility/valuation function), $v_i : [0, 1] \rightarrow \mathbb{R}^+ \cup \{0\}$, that indicates how agent i values different parts of the cake, and given a piece of cake X , the value for X is defined by $V_i(X) = \int_X v_i(x) dx = \sum_{I \in X} \int_{x \in I} v_i(x) dx$. (Note that the fact that the valuation functions are integrable implies additivity—meaning for two disjoint intervals I and I' , $V_i(I \cup I') = V_i(I) + V_i(I')$.) Without loss of generality, we assume that the valuation functions are *normalized*, i.e., $V_i([0, 1]) = 1$ for all $i \in [n]$. Additionally, we say that agent i is *hungry* if $v_i(x) > 0$ for all $x \in [0, 1]$.

2.1 Input Models

The way we have defined the valuation functions above, it is not necessary that they have a discrete representation. Therefore, most of the work in cake cutting assumes a query model—commonly known as the Robertson-Webb query model—that allows only the following two types of queries.

1. Eval query: given an interval $[x, y]$, **eval**(i, x, y) asks agent i for its value for $[x, y]$, i.e., **eval**(i, x, y) = $V_i(x, y)$.
2. Cut query: given a point x and $r \in [0, 1]$, **cut**(i, x, r) asks agent i for the minimum (i.e., the leftmost) point y such that $V_i(x, y) = r$.

While the Robertson-Webb model is the most widely studied, there is recent work (e.g., [Chen *et al.*, 2010; 2013], [Bei *et al.*, 2012], [Aziz and Ye, 2014]) which restricts the agents' valuation functions to ones that have a concise representation, thus in turn giving rise to another model (referred to as the direct-revelation model) where the entire valuation functions of the agents are given as input to a mechanism. In this paper, we consider one fundamental class of restricted valuations, namely, piecewise constant valuations. An agent i is said to have a piecewise constant valuation if $[0, 1]$ can be partitioned into a finite number of intervals such that v_i is a constant over each interval (i.e., it is a step function). A special case of piecewise constant valuations are the piecewise uniform valuations where, for some constant c , the function can only take the values 0 or c . One of the reasons why piecewise constant valuations are interesting is because they are expressive enough to be able to approximate any general valuation function.

In this paper, we make use of both input models. All our negative (impossibility) results are in the direct-revelation model, while the algorithm we present is in the Robertson-Webb model. Note that the existence of a mechanism in the Robertson-Webb model implies an existence in the direct-revelation model, but the converse is not necessarily true, and hence our negative results, in particular, are strong. Also, all our negative results are true for even the restricted class of piecewise constant valuations and hence naturally carry over to anything more general (like piecewise linear valuations).

2.2 Properties of Cake Cutting Mechanisms

A direct-revelation cake cutting mechanism \mathcal{M} takes the valuation functions (v_1, \dots, v_n) of the agents (which we refer to as a profile) as input and it outputs an allocation $A = (A_1, \dots, A_n)$, where A_i is the allocation to agent i , and $\forall i, j (\neq i)$, A_i and A_j are disjoint (except at the boundaries of the intervals). While in general A_i can be any arbitrary piece of cake, sometimes we concern ourselves only with *contiguous allocations* where each A_i is a single interval. Throughout, we assume that the mechanism allocates (to some agent) all the pieces of cake that are valued at greater than zero by at least one of the agents. Additionally, we often use $\mathcal{M}_i(v_i, v_{-i})$ to denote A_i , where v_{-i} denotes the valuation functions of all the agents except i .

A mechanism is said to be *non-wasteful* if it allocates every piece that is desired—meaning that the piece is assigned a value greater than zero—by at least one agent to an agent who desires it. More formally, if I is a subinterval of $[0, 1]$ and $D(I) = \{i \in [n] \mid V_i(I) > 0\}$, then a mechanism is non-wasteful if it always makes an allocation A such that $\forall I, I \subseteq \cup_{i \in D(I)} A_i$. In this paper, we consider both non-wasteful and wasteful mechanisms. Additionally, for all our negative results we assume *free disposal*, meaning that the mechanism can throw away pieces of cake that aren't valued (at greater than zero) by any of the agents without incurring a cost. Note that the existence of a mechanism without the free disposal assumption implies the existence of a mechanism with the free disposal assumption, but the converse is not necessarily true. Hence, in the context of negative results, such an assumption only makes them stronger.

In this paper, we use (approximate) proportionality as our main fairness criterion. In the definition below, proportionality refers to the special case when $\epsilon = 0$.

Definition 1 (ϵ -proportionality). *For any $\epsilon \in [0, \frac{1}{n})$, a mechanism \mathcal{M} is said to be ϵ -proportional if it always returns an allocation (A_1, \dots, A_n) such that $\forall i \in [n], V_i(A_i) \geq \frac{1}{n} - \epsilon$.*

Along with (approximate) proportionality, the other major property we are concerned with is (approximate) strategyproofness. Informally, for any $\epsilon \in [0, 1]$, a cake cutting mechanism \mathcal{M} is said to be ϵ -strategyproof if an agent can gain a utility of at most ϵ by misreporting her valuation function. More formally, we have the following definition.

Definition 2 (ϵ -strategyproofness). *For any $\epsilon \in [0, 1]$, a mechanism \mathcal{M} is said to be ϵ -strategyproof if for every agent $i \in [n]$, and for all v'_i, v_{-i} , $V_i(\mathcal{M}_i(v_i, v_{-i})) \geq V_i(\mathcal{M}_i(v'_i, v_{-i})) - \epsilon$.*

Strategyproofness refers to the special case in the above definition when $\epsilon = 0$.

In addition to the above mentioned properties, we also talk about (approximate) envy-freeness and Pareto-efficiency. For an $\epsilon \in [0, 1]$, a cake cutting mechanism is said to be ϵ -*envy-free* if it always returns an allocation (A_1, \dots, A_n) such that $\forall i, j \in [n], V_i(A_i) \geq V_i(A_j) - \epsilon$. Envy-freeness refers to the special case in the above definition when $\epsilon = 0$. A cake cutting mechanism is said to be *Pareto-efficient* if it always returns an allocation (A_1, \dots, A_n) that is not Pareto-dominated—meaning that there is no other allocation (A'_1, \dots, A'_n) such that $\forall i \in [n], V_i(A'_i) \geq V_i(A_i)$, with the inequality being strict for at least one agent. Note that non-wastefulness, as defined above, can be considered as a weak form of efficiency where all it is doing is restricting a mechanism from allocating a piece of zero value to an agent if there is some other agent who has a non-zero value for it.

3 Non-wasteful Mechanisms

Chen *et al.* [2010; 2013] were the first to consider direct-revelation mechanisms for cake cutting and they proposed a polynomial-time deterministic mechanism that is strategyproof, proportional, envy-free, and Pareto-efficient for piecewise uniform valuations. In light of such a result, one natural question that arises is if there exists a mechanism for at least the more general class of piecewise constant valuations that is strategyproof, Pareto-efficient, and fair (for some notion of fairness). We already have an answer to this question due to Schummer [1996] who considered the predefined divisible goods setting and showed that the only mechanism that is strategyproof and Pareto-efficient is a dictatorship. And since this predefined goods setting is a special case of the setting with piecewise constant valuations, we know that there is no hope for achieving strategyproofness, Pareto-efficiency, and any notion of fairness. Therefore, in this context, another notion to consider is non-wastefulness which, as defined in Section 2, is a weaker notion of efficiency. In fact, Aziz and Ye [2014] considered this notion and showed an impossibility result saying that there is no strategyproof, robust proportional (which is a much stronger notion of proportionality; see [Aziz and Ye, 2014]), and non-wasteful mechanism

[Aziz and Ye, 2014, Theorem 7]. Below, we show a stronger impossibility which says that there is no deterministic and non-wasteful mechanism that is even ϵ_1 -strategyproof and ϵ_2 -proportional for ϵ_1, ϵ_2 such that $0 \leq 3\epsilon_1 + \epsilon_2 < \frac{1}{n}$.

Theorem 1. *For any $n \geq 2$ and $0 \leq \epsilon_1, \epsilon_2 < \frac{1}{n}$ such that $3\epsilon_1 + \epsilon_2 < \frac{1}{n}$, there is no deterministic and non-wasteful mechanism for n agents with piecewise constant valuations that is ϵ_1 -strategyproof and ϵ_2 -proportional.*

Proof. Let us assume that there exists a deterministic, ϵ_1 -strategyproof, ϵ_2 -proportional, and non-wasteful mechanism \mathcal{M} for n agents. First, consider the profile (u, u, y, \dots, y) , where

$$u(x) = \begin{cases} \frac{n}{2} & x \in [0, \frac{2}{n}] \\ 0 & x \in (\frac{2}{n}, 1] \end{cases} \quad y(x) = \begin{cases} 0 & x \in [0, \frac{2}{n}] \\ \frac{n}{n-2} & x \in (\frac{2}{n}, 1]. \end{cases}$$

Let \mathcal{M} make an allocation A_1, A_2 to agent 1 and agent 2, respectively. Without loss of generality, we can assume that $|A_1| \geq |A_2|$. Now, since \mathcal{M} is non-wasteful, $|A_1| + |A_2| = \frac{2}{n}$, and so this implies that $|A_1| \geq \frac{1}{n}$.

Next, let us consider the profile (v, u, y, \dots, y) , where $v(x) = \frac{1}{|A_1|}$ when $x \in A_1$ and 0 otherwise. Since \mathcal{M} is non-wasteful and ϵ_1 -strategyproof, agent 1 has to get a utility of at least $1 - \epsilon_1$ from the allocation in this profile, because if it allocates anything less, then agent 1 will deviate to Profile 1. Therefore, if \mathcal{M} allocates A'_1 and A'_2 to agent 1 and 2, respectively, then $(1 - \epsilon_1) \cdot |A_1| \leq |A'_1|$ and so $|A'_2| \leq |A_2| + \epsilon_1 \cdot |A_1|$.

Finally, consider the profile (v, w, y, \dots, y) , where

$$w(x) = \begin{cases} \frac{1-\delta}{|A_1|} & x \in A_1 \\ \frac{\delta}{|A_2|} & x \in A_2 \\ 0 & \text{otherwise} \end{cases}$$

Now, here, since \mathcal{M} is non-wasteful and because agent 1 does not desire any part of A_2 , agent 2 gets at least A_2 . However, agent 2 has a total utility of only δ for A_2 , and since we assumed that \mathcal{M} is also ϵ_2 -proportional, it has to get a piece of length, say, ℓ , of A_1 , where

$$\delta + \frac{\ell}{|A_1|}(1 - \delta) \geq \frac{1}{n} - \epsilon_2 \implies \ell \geq |A_1| \cdot \frac{(\frac{1}{n} - \epsilon_2 - \delta)}{(1 - \delta)}.$$

If \mathcal{M} allocates A''_1 and A''_2 to agent 1 and 2, respectively, in this profile, then $|A''_2| \geq |A_2| + \ell \geq |A_2| + |A_1| \cdot \frac{(\frac{1}{n} - \epsilon_2 - \delta)}{(1 - \delta)}$. Also, from the profile (v, u, y, \dots, y) we know that $|A'_2| \leq |A_2| + \epsilon_1 \cdot |A_1|$. Therefore, if $\frac{n}{2} \left(|A_2| + |A_1| \cdot \frac{(\frac{1}{n} - \epsilon_2 - \delta)}{(1 - \delta)} \right) > \frac{n}{2} (|A_2| + \epsilon_1 \cdot |A_1|) + \epsilon_1$, then agent 2 can deviate from (v, u, y, \dots, y) to (v, w, y, \dots, y) and as a result gain strictly more than ϵ_1 . So, considering this inequality, we have,

$$\begin{aligned} \frac{n}{2} \left(|A_2| + |A_1| \cdot \frac{(\frac{1}{n} - \epsilon_2 - \delta)}{(1 - \delta)} \right) &> \frac{n}{2} (|A_2| + \epsilon_1 |A_1|) + \epsilon_1 \\ \frac{1}{n} \cdot \left(\frac{(\frac{1}{n} - \epsilon_2 - \delta)}{(1 - \delta)} - \epsilon_1 \right) &> \frac{2}{n} \epsilon_1 \quad (\text{as } |A_1| \geq \frac{1}{n}) \\ \delta &< \frac{1 - n(3\epsilon_1 + \epsilon_2)}{n(1 - 3\epsilon_1)}. \end{aligned}$$

This in turn implies that for every $0 \leq \epsilon_1, \epsilon_2 < \frac{1}{n}$ such that $3\epsilon_1 + \epsilon_2 < \frac{1}{n}$, we can always find a $0 < \delta < \frac{1 - n(3\epsilon_1 + \epsilon_2)}{n(1 - 3\epsilon_1)}$ to fit in the definition of the function $w(x)$ such that the inequality—and hence our theorem—will be true. \square

With respect to the above theorem, we highlight the following corollaries which are basically special cases when ϵ_1 and ϵ_2 are zero, respectively. The first one rules out the possibility of guaranteeing some value (however small) to all the agents if the mechanism has to be non-wasteful and strategyproof. The second one, on the other hand, provides a lower bound for the natural question on whether we can come up with mechanisms that can guarantee proportionality and non-wastefulness and at the same time allow only a maximum gain of ϵ , where $0 \leq \epsilon < (1 - \frac{1}{n})$, from misreporting. Unfortunately, we do not have an accompanying upper bound result for this and so the question of an upper bound remains open.

Corollary 2. *For any $n \geq 2$ and $0 \leq \epsilon < \frac{1}{n}$, there is no deterministic, strategyproof, ϵ -proportional, and non-wasteful mechanism for n agents with piecewise constant valuations.*

Corollary 3. *For any $n \geq 2$ and $0 \leq \epsilon < \frac{1}{3n}$, there is no deterministic, ϵ -strategyproof, proportional, and non-wasteful mechanism for n agents with piecewise constant valuations.*

4 Removing Non-wastefulness

While non-wastefulness is certainly a desirable property, it seems like the negative results in the previous section were largely driven by this constraint as it severely restricts the kind of allocations that a mechanism can make. So, what if we remove this constraint? Do similar impossibilities still hold, or are there mechanisms that satisfy some of those properties? These are the questions we try to answer here.

Chen *et al.* [2010; 2013] provided a deterministic and strategyproof mechanism for piecewise uniform valuations and one of the main open questions they had posed was on generalizing their mechanism for piecewise constant valuations. Subsequently, Aziz and Ye [2014] and Brânzei and Miltersen [2015] had posed the same question regarding the existence of mechanisms that are strategyproof and proportional for piecewise constant valuations. Here we address a special case of this question when the mechanism makes contiguous allocations and show that there exists no deterministic direct-revelation mechanism that is strategyproof and ϵ -proportional, for any $0 \leq \epsilon < \frac{1}{n}$. Additionally, we also prove a stronger result (Proposition 4) for the case of two agents. Informally, at a high-level, the proofs of both of these results follow the same theme as in the proof of Theorem 1 where we construct valuation profiles with an aim of arriving at a contradiction. However, unlike in Theorem 1, we now no longer have the non-wasteful constraint, and so this requires some additional arguments which use the fact that the allocations are contiguous. Due to space constraints, the proofs of both of the statements below have been deferred to the full version [Menon and Larson, 2017].

Proposition 4. *For any $0 \leq \epsilon_1, \epsilon_2 < \frac{1}{2}$ such that $\epsilon_1 + \epsilon_2 < \frac{1}{2}$, there exists no deterministic mechanism for two hungry agents with piecewise constant valuations that*

makes contiguous allocations and is ϵ_1 -strategyproof and ϵ_2 -proportional.

As an aside, note that the $\epsilon_2 = 0$ case in Proposition 4 essentially says that if we insist on contiguous and proportional allocations, then one cannot provide any guarantee on strategyproofness (i.e., such mechanisms cannot be ϵ -strategyproof for any feasible ϵ). Later on, in Theorem 8, we show how one can circumvent this and obtain a $\frac{1}{4}$ -strategyproof and proportional mechanism for two agents by giving each of them at most two contiguous pieces.

Theorem 5. *For any $n \geq 2$ and $0 \leq \epsilon < \frac{1}{n}$, there is no deterministic mechanism for n agents with piecewise constant valuations that makes contiguous allocations and is strategyproof and ϵ -proportional.*

Given Theorem 5, the first natural question that arises is if we can extend it to the case when the mechanism can make arbitrary allocations. Unfortunately, we do not have an answer to this question. Instead, below, we show a connection between direct-revelation mechanisms and mechanisms in the Robertson-Webb model for the case when agents have piecewise constant valuations. We believe that it will be useful in either of the cases—i.e., if the answer to the above question is either in the positive or if we want to prove that such a mechanism does not exist. In the case that such a mechanism exists, the connection basically shows that when given an upper bound on the maximum number of breakpoints in any agent’s valuation function, one can construct a mechanism in the Robertson-Webb model that approximately achieves both of the properties. And in the case that one wants to prove that such a mechanism does not exist, one can use this connection to prove a non-existence (of mechanisms that are ϵ -strategyproof and ϵ -proportional for some $\epsilon > 0$) in the Robertson-Webb model and then map it back to show that there is no (finite) direct-revelation mechanism that is strategyproof and proportional.

Proposition 6. *Let P_1, P_2, P_3 denote the property of strategyproofness, proportionality, and envy-freeness, respectively, and let $P \subseteq \{P_1, P_2, P_3\}$. Given k , an upper bound on the maximum number of breakpoints in any agent’s valuation function, and any $n \geq 2$, if there exists a (finite) deterministic direct-revelation mechanism for n agents with piecewise constant valuations that satisfies $P_i, \forall P_i \in P$, then, for any $\epsilon > 0$, there exists a mechanism in the Robertson-Webb model for n agents with piecewise constant valuations that asks at most $(n \lfloor \frac{2k}{\epsilon} \rfloor)$ queries on each input and is ϵ - $P_i, \forall P_i \in P$.*

Proof (sketch). Since in the Robertson-Webb model we only have query access to the valuation functions, the first step is to basically “learn” these functions to a sufficiently good approximation using only **eval** and **cut** queries. Subsequently, once we have that, we can essentially feed these functions to the direct-revelation mechanism in order to create a mechanism in the Robertson-Webb model that achieves an approximate version of all the properties. For the first step, the key observation is following claim.

Claim 1. *Given an $\epsilon > 0$, we can use $\lfloor \frac{2k}{\epsilon} \rfloor$ **cut** queries to ϵ -approximate any piecewise constant function $v(x)$ that has*

Input: An agent i ’s piecewise constant valuation function $v_i(x)$, $\epsilon > 0$, and k
Output: A piecewise constant function $w_i(x)$ that ϵ -approximates $v_i(x)$
1: $N \leftarrow \lfloor \frac{2k}{\epsilon} \rfloor, x_0 \leftarrow 0, x_{N+1} \leftarrow 1$
2: **for** $j \in \{1, \dots, N\}$ **do**
3: $x_j \leftarrow \text{cut}(i, x_{j-1}, \frac{\epsilon}{2k}), c_j \leftarrow \frac{\epsilon}{2k \cdot (x_j - x_{j-1})}$
4: $w_i(x) = c_j, x \in [x_{j-1}, x_j]$
5: **end for**
6: **if** $X_N \neq 1$ **then**
7: $w_i(x) = \frac{(1 - \frac{\epsilon}{2k}) \cdot N}{x_{N+1} - x_N}, x \in [x_N, x_{N+1}]$
8: **end if**
9: **return** $w_i(x)$

Algorithm 1: ϵ -approximation of a piecewise constant function using $\lfloor \frac{2k}{\epsilon} \rfloor$ **cut** queries.

at most k breakpoints by another piecewise constant function $w(x)$ such that $W([0, 1]) = 1$ and for any piece of cake X ,

$$\int_X v(x) dx - \frac{\epsilon}{2} \leq \int_X w(x) dx \leq \int_X v(x) dx + \frac{\epsilon}{2}. \quad (1)$$

Proof (sketch). Algorithm 1 describes how to convert $v(x)$ to $w(x)$ by using $\lfloor \frac{2k}{\epsilon} \rfloor$ **cut** queries. It is easy to see that $W([0, 1]) = 1$. To prove that for any piece of cake X , $\int_X v(x) dx - \frac{\epsilon}{2} \leq \int_X w(x) dx \leq \int_X v(x) dx + \frac{\epsilon}{2}$, consider some arbitrary X . If we denote the breakpoints of v by $\{b_1, \dots, b_k\}$ (since there are only a maximum of k breakpoints), where $0 = b_0 < b_1 < \dots < b_{k+1} = 1$, and if we let $C_i = X \cap [x_{i-1}, x_i]$, then we know that $\int_X w(x) dx = \sum_{i=1}^{N+1} \int_{C_i} w(x) dx$. Additionally, also note that if $[x_{i-1}, x_i]$ does not contain any breakpoint b_j of the original function $v(x)$, then $\int_{C_i} w(x) dx = \int_{C_i} v(x) dx$ since we know that $v(x)$ is constant over this interval (as there are no breakpoints) and so $v(x) = \frac{\epsilon}{2k \cdot (x_i - x_{i-1})} = w(x), \forall x \in [x_{i-1}, x_i]$. This, along with the fact that $v(x)$ has at most k breakpoints, implies that there are at most k C_i ’s for which $\int_{C_i} w(x) dx \neq \int_{C_i} v(x) dx$. Now, if we assume without loss of generality that these are $C_1 \dots, C_k$, then using the fact that $\int_{C_i} w(x) dx \leq \frac{\epsilon}{2k}$ and $\int_{C_i} v(x) dx \leq \frac{\epsilon}{2k}$ we have that for every $i \in \{1, \dots, k\}$,

$$-\frac{\epsilon}{2k} \leq \int_{C_i} w(x) dx - \int_{C_i} v(x) dx \leq \frac{\epsilon}{2k}. \quad (2)$$

Now, since $\int_X [w(x) - v(x)] dx = \sum_{i=1}^k (\int_{C_i} [w(x) - v(x)] dx)$, we can now use equation 2 to prove our claim. \square

Given the above claim, we can now build a mechanism \mathcal{M}^{RW} in the Robertson-Webb model in the following way, where (v_1, \dots, v_n) are the original valuation functions and (w_1, \dots, w_n) are the corresponding ϵ -approximated functions.

$$\forall i \in [n], \mathcal{M}_i^{RW}(v_1, \dots, v_n) = \mathcal{M}_i(w_1, \dots, w_n). \quad (3)$$

It is clear that \mathcal{M}^{RW} makes at most $(n \lfloor \frac{2k}{\epsilon} \rfloor)$ queries on each input. To prove that \mathcal{M}^{RW} is ϵ - P_i for all $P_i \in P$ satisfied by \mathcal{M} , let us consider each of the properties separately.

i) When \mathcal{M} satisfies P_1 (strategyproofness). Since \mathcal{M} is strategyproof, the only way for an agent i with a valuation function v_i to manipulate \mathcal{M}^{RW} is to pretend as if its function is some v'_i and answer the queries accordingly. This in turn will result in w'_i being the ϵ -approximated function instead of w_i if it reported truthfully. However, we know that $\forall w'_i, w_{-i}, W_i(\mathcal{M}_i(w_i, w_{-i})) \geq W_i(\mathcal{M}_i(w'_i, w_{-i}))$. This in turn implies that using equation 1 we have,

$$V_i(\mathcal{M}_i(w_i, w_{-i})) + \frac{\epsilon}{2} \geq W_i(\mathcal{M}_i(w_i, w_{-i})) \geq W_i(\mathcal{M}_i(w'_i, w_{-i})) \geq V_i(\mathcal{M}_i(w'_i, w_{-i})) - \frac{\epsilon}{2}.$$

And so, now, using equation 3, we have that for all v'_i, v_{-i} , $V_i(\mathcal{M}_i^{RW}(v'_i, v_{-i})) - V_i(\mathcal{M}_i^{RW}(v_i, v_{-i})) \leq \epsilon$.

The other two cases can be proved similarly and we defer the complete proof to the full version. \square

Although, as mentioned above, it remains open as to whether a strategyproof and proportional mechanism exists, another goal one could still consider is on finding better proportional mechanisms (in terms of ϵ -strategyproofness). Below, we begin the pursuit of this goal and as our final question in this paper ask if there are positive results at least regarding mechanisms that are proportional and ϵ -strategyproof for some $\epsilon < (1 - \frac{1}{n})$. We answer it with a “Yes” and in particular we first prove—by showing a bound on the maximum gain an agent can get by misreporting—that the well-known Even-Paz algorithm [Even and Paz, 1984] satisfies the above criterion for the case when there are at least four agents. And following this, we present a minor modification to the same that has better incentive compatible properties than the original Even-Paz algorithm for the cases when there are two or three agents. Due to space constraints, proofs of Proposition 7 and Theorem 8 have been deferred to the full version.

Proposition 7. *For $n \geq 2$ agents, the Even-Paz algorithm is deterministic, proportional, and ϵ -strategyproof, where*

$$\epsilon = \begin{cases} \frac{1}{2} & \text{when } n = 2 \text{ or } n = 4 \\ \frac{2}{3} & \text{when } n = 3 \text{ or } n = 5 \\ 1 - \frac{2}{n} & \text{when } n \geq 6. \end{cases}$$

Next, we show that the bounds on ϵ can be improved for the cases when there are two or three agents. In particular, we show that Algorithm 2 is proportional and ϵ -strategyproof, where $\epsilon = (1 - \frac{3}{2n})$ or $\epsilon = (1 - \frac{3}{2n} + \frac{1}{2n^2})$ depending on whether n is even or odd, respectively. And so, for the cases when there are two or three agents, the gain is $\frac{1}{4}$ or $\frac{5}{9}$, respectively, as opposed to $\frac{1}{2}$ or $\frac{2}{3}$ in the Even-Paz algorithm.

Theorem 8. *Algorithm 2 is deterministic, proportional, and ϵ -strategyproof, where*

$$\epsilon = \begin{cases} 1 - \frac{3}{2n} & \text{when } n \text{ is even} \\ 1 - \frac{3}{2n} + \frac{1}{2n^2} & \text{when } n \text{ is odd.} \end{cases}$$

5 Discussion

One of the main open questions raised by Chen *et al.* [2010; 2013] (and also by Aziz and Ye [2014] and Brânzei and Miltersen [2015]) was on the existence of deterministic mechanisms that are strategyproof and fair for piecewise constant

Procedure: modifiedEP($[a, b], S$)
Input: ($[a, b], S$), where $[a, b] \subseteq [0, 1]$ is the cake to be proportionally allocated among $k = |S|$ agents
1: If $k == 1$, then **return** ($[a, b]$)
2: **for** each $i \in S$ **do**
3: $c_i = \text{cut}\left(i, a, \frac{\lfloor \frac{k}{2} \rfloor}{k} \cdot V_i(a, b)\right)$
4: **end for**
5: Let $\{d_1, \dots, d_k\}$ be the sorted order of the c_i 's, $S_L \leftarrow \{i \in S \mid c_i \leq d_{\lfloor \frac{k}{2} \rfloor}\}$, and $S_R \leftarrow S \setminus S_L$.
6: Proportionally allocate the cake $[d_{\lfloor \frac{k}{2} \rfloor}, d_{\lfloor \frac{k}{2} \rfloor + 1}]$ among the k agents using the Even-Paz algorithm.
7: Recursively call modifiedEP on $[a, d_{\lfloor \frac{k}{2} \rfloor}]$ and $[d_{\lfloor \frac{k}{2} \rfloor + 1}, b]$ with S_L and S_R , respectively. Return the allocations so formed along with the proportional allocation of the piece $[d_{\lfloor \frac{k}{2} \rfloor}, d_{\lfloor \frac{k}{2} \rfloor + 1}]$.
Main:
8: output modifiedEP($[0, 1], S = \{1, \dots, n\}$).

Algorithm 2: Modified Even-Paz algorithm.

valuations. Here we addressed a special case of this and we showed (in Theorem 5) that for any $n \geq 2$ agents there is no deterministic mechanism that makes contiguous allocations and is strategyproof and even ϵ -proportional. Although the contiguous allocations constraint captures some interesting scenarios, we believe that answering the above question without this (rather strong) constraint is important. In particular, when given the fact that there exists randomized mechanisms that are a) truthful in expectation, proportional, and envy-free for piecewise linear valuations in the direct-revelation model [Chen *et al.*, 2013] and b) truthful in expectation and ϵ -proportional in the Robertson-Webb model [Brânzei and Miltersen, 2015], we believe that resolving this question can give us a clearer picture regarding the limits of determinism.

Moving away from the above result, in Section 4 we showed (in Theorem 8) that there exists a proportional mechanism that has better incentive-compatible properties than the Even-Paz algorithm for the cases when there are two or three agents. With regards to this result, one of the main questions that arise is the following and we suggest it as an open problem: for $n \geq 2$ agents, what is the minimum achievable ϵ for which there exists a deterministic mechanism that is proportional and ϵ -strategyproof? It is unclear how we can provide a lower bound, or improve the proposed mechanism to get smaller values of ϵ . And also, at the same time there is the question of whether we can make Algorithm 2 non-wasteful when agents have concisely representable valuation functions. Therefore, given our lower bound (Theorem 1) in Section 3, another problem that remains open is whether we can guarantee non-wastefulness along with proportionality and ϵ -strategyproofness, where $\frac{1}{3n} \leq \epsilon < (1 - \frac{1}{n})$.

Acknowledgments

We thank Simina Brânzei for useful discussions.

References

- [Alijani *et al.*, 2017] Reza Alijani, Majid Farhadi, Mohammad Ghodsi, Masoud Seddighin, and Ahmad S Tajik. Envy-free mechanisms with minimum number of cuts. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 312–318, 2017.
- [Aziz and Mackenzie, 2016] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 416–427, 2016.
- [Aziz and Ye, 2014] Haris Aziz and Chun Ye. Cake cutting algorithms for piecewise constant and piecewise uniform valuations. In *Proceedings of the 10th International Conference on Web and Internet Economics (WINE)*, pages 1–14. Springer, 2014.
- [Bei *et al.*, 2012] Xiaohui Bei, Ning Chen, Xia Hua, Biaoshuai Tao, and Endong Yang. Optimal proportional cake cutting with connected pieces. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1263–1269, 2012.
- [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 *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 482–488, 2015.
- [Chen *et al.*, 2010] Yiling Chen, John K Lai, David Parkes, and Ariel D Procaccia. Truth, justice, and cake cutting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, pages 626–631, 2010.
- [Chen *et al.*, 2013] Yiling Chen, John K Lai, David C Parkes, and Ariel D Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1):284–297, 2013.
- [Cheung, 2016] Yun Kuen Cheung. Better strategyproof mechanisms without payments or prior - an analytic approach. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 194–200, 2016.
- [Chevalleyre *et al.*, 2006] Yann Chevalleyre, Paul E Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaitre, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A Rodriguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica*, 30(1):3–31, 2006.
- [Cole *et al.*, 2013] Richard Cole, Vasilis Gkatzelis, and Gagan Goel. Mechanism design for fair division: allocating divisible items without payments. In *ACM Conference on Electronic Commerce (EC)*, pages 251–268, 2013.
- [Even and Paz, 1984] Shimon Even and Azaria Paz. A note on cake cutting. *Discrete Applied Mathematics*, 7(3):285–296, 1984.
- [Kurokawa *et al.*, 2013] David Kurokawa, John K Lai, and Ariel D Procaccia. How to cut a cake before the party ends. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, pages 555–561, 2013.
- [Li *et al.*, 2015] Minming Li, Jialin Zhang, and Qiang Zhang. Truthful cake cutting mechanisms with externalities: Do not make them care for others too much. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 589–595, 2015.
- [Maya and Nisan, 2012] Avishay Maya and Noam Nisan. Incentive compatible two player cake cutting. In *Proceedings of the 8th Workshop on Internet and Network Economics (WINE)*, pages 170–183. Springer, 2012.
- [Menon and Larson, 2017] Vijay Menon and Kate Larson. Deterministic, strategyproof, and fair cake cutting. *arXiv preprint*, abs/1705.06306, 2017.
- [Mossel and Tamuz, 2010] Elchanan Mossel and Omer Tamuz. Truthful fair division. In *International Symposium on Algorithmic Game Theory (SAGT)*, pages 288–299. Springer, 2010.
- [Procaccia, 2016] Ariel D Procaccia. Cake cutting algorithms. In Vincent Conitzer, Ulle Endriss, Jerome Lang, and Ariel D Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13. Cambridge University Press, 2016.
- [Robertson and Webb, 1998] Jack Robertson and William Webb. *Cake-Cutting Algorithms: Be Fair If You Can*. A.K. Peters., 1998.
- [Schummer, 1996] James Schummer. Strategy-proofness versus efficiency on restricted domains of exchange economies. *Social Choice and Welfare*, 14(1):47–56, 1996.
- [Steinhaus, 1948] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.