Truth, Justice, and Cake Cutting[1]

Yiling Chen, John K. Lai, David C. Parkes, and Ariel D. Procaccia

Superman: "I'm here to fight for truth, justice, and the American way."
Lois Lane: "You're gonna wind up fighting every elected official in this country!"

Superman (1978)

Abstract

Cake cutting is a common metaphor for the division of a heterogeneous divisible good. There are numerous papers that study the problem of fairly dividing a cake; a small number of them also take into account self-interested agents and consequent strategic issues, but these papers focus on fairness and consider a strikingly weak notion of truthfulness. In this paper we investigate the problem of cutting a cake in a way that is truthful and fair, where for the first time our notion of dominant strategy truthfulness is the ubiquitous one in social choice and computer science. We design both deterministic and randomized cake cutting algorithms that are truthful and fair under different assumptions with respect to the valuation functions of the agents.

1 Introduction

The need for resource allocation arises in many AI domains, and in particular in multiagent systems. This has led to a wide interest in the field known as *Multiagent Resource Allocation*, and to various applications of resource allocation techniques (see the survey by Chevalyere et al. [7]). Resource allocation problems deal with either *divisible* or *indivisible* resources, where the distinction is based on whether any fraction of a resource can be given to an agent.

Cutting a cake is often used as a metaphor for allocating a divisible good. The difficulty is not cutting the cake into pieces of equal size, but rather that the cake is not uniformly tasty: different agents prefer different parts of the cake, depending, e.g., on whether the toppings are strawberries or cookies. The goal is to divide the cake in a way that is "fair"; the definition of fairness is a nontrivial issue in itself, which we discuss in the sequel. The cake cutting problem dates back to the 1940s, and for over sixty years has attracted the attention of mathematicians, economists, and political scientists. While most of the work in artificial intelligence, and computer science in general, has focused on the allocation of indivisible resources, recent years have seen an increasing interest among computer scientists in the allocation of divisible resources (see, e.g, [9, 10, 15]).

Slightly more formally, the cake is represented by the interval $[0, 1]$. Each of n agents has a valuation function over the cake, which assigns a value to every given piece of cake and is additive. The goal is to find a partition of the cake among the agents (while possibly throwing a piece away) that satisfies one or several fairness criteria. In this paper we consider the two most prominent criteria. A *proportional* allocation is one where the value each agent has for its own piece of cake is at least $1/n$ of the value it assigns to the entire cake. An

[1] A version that is similar to this extended abstract will appear in the proceedings of AAAI'10. The full version of the paper, which includes all omitted proofs and a longer exposition, will shortly be available online. The paper was presented in the Harvard EconCS seminar (February 2010) and in a workshop on prior-free mechanism design in Guanajuato, Mexico (May 2010).

envy-free (EF) allocation is one where the value each agent assigns to its own piece of cake is at least as high as the value it assigns to any other agent's piece of cake. There is a rather large body of literature on fairly cutting a cake according to these two criteria (see, e.g., the books by Robertson and Webb [16] and Brams and Taylor [6]).

So far we have briefly discussed "justice", but have not yet mentioned "truth." Taking the game-theoretic point of view, an agent's valuation function is its private information, which is reported to a cake cutting algorithm. We would like an algorithm to be *truthful*, in the sense that agents are motivated to report their true valuation functions. Like fairness, this idea of truthfulness also lends itself to many interpretations. One variation, referred to as *strategy-proofness* in previous papers by Brams et al. [4, 5], assumes that an agent would report its truthful valuation rather than lie if there *exist* valuations of the other agents such that reporting truthfully yields at least as much value as lying. In the words of Brams et al., "...the players are risk-averse and never strategically announce false measures if it does not guarantee them more-valued pieces. ... Hence, a procedure is strategy-proof if no player has a strategy that dominates his true value function." [5, page 362].

The foregoing notion is strikingly weak compared to the notion of truthfulness that is common in the social choice literature. Indeed, strategy-proofness is usually taken to mean that an agent can *never* benefit by lying, that is, *for all* valuations of the other agents reporting truthfully yields at least as much value as lying. Put another way, truth-telling is a dominant strategy. This notion is worst-case, in the sense that an agent cannot benefit by lying even if it is fully knowledgeable of the valuations of the other agents. It is also the predominant one in the computer science literature, and in particular in the algorithmic mechanism design literature [14]. In order to prevent confusion we will avoid using the term "strategy-proof," and instead refer to the former notion of Brams et al. as "weak truthfulness" and to the latter standard notion as "truthfulness."

To illustrate the difference between the two notions, consider the most basic cake cutting algorithm for the case of two agents, the *Cut and Choose* algorithm.[2] Agent 1 cuts the cake into two pieces that are of equal value according to its valuation; agent 2 then chooses the piece that it prefers, giving the other piece to agent 1. This algorithm is trivially proportional and EF.[3] It is also weakly truthful, as if agent 1 divides the cake into two pieces that are unequal according to its valuation then agent 2 may prefer the piece that is worth more to agent 1. Agent 2 clearly cannot benefit by lying. However, the algorithm is not truthful. Indeed, consider the case where agent 1 would simply like to receive as much cake as possible, whereas the single-minded agent 2 is only interested in the interval $[0, \epsilon]$ where ϵ is small (for example, it may only be interested in the cherry). If agent 1 follows the protocol it would only receive half of the cake. Agent 1 can do better by reporting that it values the intervals $[0, \epsilon]$ and $[\epsilon, 1]$ equally, since then it would end up with almost the entire cake by choosing to cut pieces $[0, \epsilon], [\epsilon, 1]$.

In this paper we consider the design of truthful and fair cake cutting algorithms. To the best of our knowledge we are the first to do so. However, there is a major obstacle that must be circumvented: regardless of strategic issues, and when there are more than four agents, even finding a proportional and EF allocation in a bounded number of steps with a deterministic algorithm is a long-standing open problem! See [15] for an up-to-date discussion.[4] We shall therefore restrict ourselves to specific classes of valuation functions where efficiently finding fair allocations is a non-issue; the richness of our problem stems from our desire to additionally achieve truthfulness.

[2]This algorithm is described here with the agents taking actions; equivalently, the algorithm acts on behalf of agents using the reported valuations.

[3]Proportionality and envy-freeness coincide if there are two agents and the entire cake is allocated.

[4]To be precise, previous algorithmic work assumed that the entire cake has to be allocated, but this does not seem to be a significant restriction in the context of fairness.

Our results. We first consider deterministic algorithms. We restrict ourselves to the case where the agents hold *piecewise uniform* valuation functions, that is, each agent is interested in a collection of subintervals of $[0, 1]$ with the same marginal value for each fractional piece in each subinterval. This is the case when some parts of the cake satisfy a certain property and an agent desires as much of these parts as possible. Our main result is a deterministic algorithm for any number of agents that is truthful, proportional, EF, and polynomial-time. The proof requires many ingredients, including an application of the classic Max-Flow Min-Cut Theorem.

We next consider randomized algorithms. We slightly relax truthfulness by asking that the algorithm be *truthful in expectation*, that is, an agent cannot hope to increase its expected value by lying for any reports of other agents. For general valuations, we present a simple randomized algorithm that is truthful in expectation, and always outputs an allocation that is proportional and EF. We further establish that this algorithm is tractable under the relatively weak assumption that the agents hold *piecewise linear* valuation functions, that is where the marginal value in each subinterval of interest is a linear function.

Related work. We have recently learned of an independent working paper by Mossel and Tamuz that asks similar questions about truthful and fair cake cutting [13], but they focus on existence theorems. In particular, under general assumptions they show that there exists a mechanism that is truthful in expectation and guarantees each agent a value of more than $1/n$ in expectation. The results are then extended to the case of indivisible goods. The technical overlap between the two papers is very small; we refer the reader's attention to this overlap in a footnote in Section 4.

Thomson [17] showed that in general a truthful and Pareto-optimal algorithm must be dictatorial in the slightly different setting of pie-cutting. Note that Pareto-optimality is not a fairness property and neither implies, nor is implied by, envy-freeness or proportionality.

Our deterministic algorithm is related to a method proposed by Bogomolnaia and Moulin [3] in the context of the random assignment problem, and the network flow techniques we employ in our analysis generalize the reinterpretation of this method in terms of network flow due to Katta and Sethuraman [11]. We elaborate in Section 3.

2 Preliminaries

We consider a heterogeneous cake, represented by the interval $[0, 1]$. A *piece of cake* is a *finite* union of subintervals of $[0, 1]$. We sometimes abuse this terminology by treating a piece of cake as the set of the (inclusion-maximal) intervals that it contains. The length of the interval $I = [x, y]$, denoted $\text{len}(I)$, is $y - x$. For a piece of cake X we denote $\text{len}(X) = \sum_{I \in X} \text{len}(I)$.

The set of agents is denoted $N = \{1, \ldots, n\}$. Each agent $i \in N$ holds a private valuation function V_i, which maps given pieces of cake to the value agent i assigns them. Formally, each agent i has a *value density function*, $v_i : [0, 1] \to [0, \infty)$, that is piecewise continuous. The function v_i characterizes how agent i assigns value to different parts of the cake. The value of a piece of cake X to agent i is then defined as $V_i(X) = \int_X v_i(x)dx = \sum_{I \in X} \int_I v_i(x)dx$. We note that the valuation functions are *additive*, i.e. for any two disjoint pieces X and Y, $V_i(X \cup Y) = V_i(X) + V_i(Y)$, and *non-atomic*, that is $V_i([x, x]) = 0$ for every $x \in [0, 1]$. The last property implies that we do not have to worry about the boundaries of intervals, i.e., open and closed intervals are identical for our purposes. We further assume that the valuation functions are *normalized*, i.e. $V_i([0, 1]) = \int_0^1 v_i(x)dx = 1$.

A cake cutting algorithm is a function f from the valuation function of each agent to an allocation $(A_1, \ldots, A_n)$ of the cake such that the pieces are pairwise disjoint. For each $i \in N$ the piece A_i is allocated to agent i, and the rest of the cake, i.e., $[0, 1] \setminus \bigcup_{i \in N} A_i$, is thrown away. Here we are assuming *free disposal*, that is, the algorithm can throw away

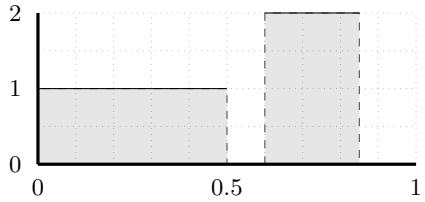

(a) Value density function for a piecewise constant valuation that is not piecewise uniform.

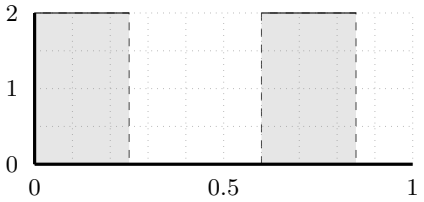

(b) Value density function for a piecewise uniform valuation.

Figure 1: An illustration of special value density functions.

resources without incurring a cost.

We say that an allocation $A_1, \ldots, A_n$ is *proportional* if for every $i \in N$, $V_i(A_i) \geq 1/n$, that is, each agent receives at least a $(1/n)$-fraction of the cake according to its own valuation. We say that an allocation is *envy-free (EF)* if for every $i, j \in N$, $V_i(A_i) \geq V_i(A_j)$, i.e., each agent prefers its own piece of cake to the piece of cake allocated to any other agent. A proportional (resp., EF) cake cutting algorithm always returns a proportional (resp., EF) allocation.

Note that when $n = 2$ proportionality implies envy-freeness. Indeed, $V_i(A_i)+V_i(A_{3-i}) \leq 1$, and hence if $V_i(A_i) \geq 1/2$ then $V_i(A_{3-i}) \leq 1/2$. Under the free disposal assumption the converse is not true. For example, an allocation that throws away the entire cake is EF but not proportional. In general, when $n > 2$ proportionality neither implies nor is implied by envy-freeness.[5]

A cake cutting algorithm f is *truthful* if when an agent lies it is allocated a piece of cake that is worth, according to its real valuation, no more than the piece of cake it was allocated when reporting truthfully. Formally, denote $A_i = f_i(V_1, \ldots, V_n)$, and let $\mathcal{V}$ be a class of valuation functions. The algorithm f is truthful if for every agent i, every collection of valuations functions $V_1, \ldots, V_n \in \mathcal{V}$, and every $V_i' \in \mathcal{V}$, it holds that $V_i(f_i(V_1, \ldots, V_n)) \geq V_i(f_i(V_1, \ldots, V_{i-1}, V_i', V_{i+1}, \ldots, V_n))$.

3 Deterministic Algorithms and Piecewise Uniform Valuations

As noted in the introduction, in general there are no known bounded deterministic proportional and EF cake cutting algorithms for more than four agents, even if one is not concerned about strategic issues. Therefore, in this section we restrict ourselves to a specific class of valuation functions.

We say that a valuation function V_i is *piecewise constant* if and only if its corresponding value density function v_i is piecewise constant, that is $[0, 1]$ can be partitioned into a finite number of intervals such that v_i is constant on each interval (see Figure 1(a)). We say that V_i is *piecewise uniform* if moreover v_i is either some constant $c \in \mathbb{R}_+$ (the same one across intervals) or zero. See Figure 1(b) for an illustration.

Piecewise uniform valuation functions imply that agent $i \in N$ is uniformly interested in a finite union of intervals, which we call its *reference piece of cake* and denote by U_i. For example, in Figure 1(b), $U_i = [0, 0.25] \cup [0.6, 0.85]$. Given a piece of cake X, it holds that $V_i(X) = \text{len}(X \cap U_i)/\text{len}(U_i)$. From the computational perspective, the size of the input to the cake cutting algorithm is the number of bits that define the boundaries of the intervals in the agents' reference pieces of cake.

[5] If free disposal is not assumed, that is, the entire cake is allocated, then envy-freeness implies proportionality for any n.

In the rest of this section we assume that the valuation functions are piecewise uniform. We believe that piecewise uniform valuations are very natural. An agent would have such a valuation function if it is simply interested in pieces of the good that have a certain property, e.g., a child only likes portions of the cake that have chocolate toppings, and wants as much cake with chocolate toppings as possible. We consider more general valuations in the next section on randomized algorithms.

3.1 A deterministic algorithm

Before introducing our algorithm we present some required notation. Let $S \subseteq N$ be a subset of agents and let X be a piece of cake. Let $D(S, X)$ denote the portions of X that are valued by at least one agent in S. Formally, $D(S, X) = \left(\bigcup_{i \in S} U_i\right) \cap X$, and is itself a union of intervals.

Let $\text{avg}(S, X) = \text{len}(D(S, X))/|S|$ denote the average length of intervals in X desired by at least one agent in S. We say that an allocation is *exact* with respect to S and X if it allocates to each agent in S a piece of cake of length $\text{avg}(S, X)$ comprised *only* of desired intervals. Clearly this requires allocating all of $D(S, X)$ since the total length of allocated intervals is $\text{avg}(S, X) \cdot |S| = \text{len}(D(S, X))$. Suppose $S = \{1, 2\}$ and $X = [0, 1]$: if $U_1 = U_2 = [0, 0.2]$ then agents 1 and 2 receiving $[0, 0.1]$ and $[0.1, 0.2]$ respectively is an exact allocation; but if $U_1 = [0, 0.2], U_2 = [0.3, 0.7]$ then there is no exact allocation.

The deterministic algorithm for n agents with piecewise uniform valuations is a recursive algorithm that finds a subset of agents with a certain property, makes the allocation decision for that subset, and then makes a recursive call on the remaining agents and the remaining intervals. Specifically, for a given set of agents $S \subseteq N$ and a remaining piece of cake to be allocated X, we find the subset $S' \subseteq S$ of agents with the smallest $\text{avg}(S', X)$. We then give an exact allocation of $D(S', X)$ to S'. We recurse on $S \setminus S'$ and the intervals not desired by any agent in S', i.e. $X \setminus D(S', X)$. The pseudocode of the algorithm is given as Algorithm 1.

Algorithm 1 $(V_1, \ldots, V_n)$

1. SubRoutine$(\{1, \ldots, n\}, [0, 1], (V_1, \ldots, V_n))$

SubRoutine$(S,\ X,\ V_1, \ldots, V_n)$:

1. If $S = \emptyset$, return.
2. Let $S_{\min} \in \underset{S' \subseteq S}{\text{argmin}}\ \text{avg}(S', X)$ (breaking ties arbitrarily).
3. Let $E_1, \ldots, E_n$ be an exact allocation with respect to $S_{\min}, X$ (breaking ties arbitrarily). For each $i \in S_{\min}$, set $A_i = E_i$.
4. SubRoutine$(S \setminus S_{\min}, X \setminus D(S_{\min}, X), (V_1, \ldots, V_n))$.

In particular, Steps 2 and 3 of SubRoutine imply that if $S = \{i\}$ then $A_i = D(S, X)$. For example, suppose $X = [0, 1]$, $U_1 = [0, 0.1]$, $U_2 = [0, 0.39]$, and $U_3 = [0, 0.6]$. In this case, the subset with the smallest average is $\{1\}$, so agent 1 receives all of $[0, 0.1]$ and we recurse on $\{2, 3\}, [0.1, 1]$. In the recursive call, set $\{2\}$ has average 0.39 - 0.1 = 0.29, set $\{3\}$ has average 0.6 - 0.1 = 0.5, and set $\{2, 3\}$ has average $(0.6 - 0.1)/2 = 0.25$. As a result, the entire set $\{2, 3\}$ is chosen as the set with smallest average, and an exact allocation of $[0.1, 1.0]$ is given to agents 2 and 3. One possible allocation is to give agent 2 $[0.1, 0.35]$ and agent 3 $[0.35, 0.6]$. Note that, if agent 1 uniformly values $[0, 0.2]$ instead, the first call would choose $\{1, 2\}$ as the subset with the smallest average, equally allocating $[0, 0.39]$ between agents 1 and 2 and giving the rest, $[0.39, 0.6]$, to agent 3.

An analysis of the two agent algorithm. To gain intuition, consider the case of two

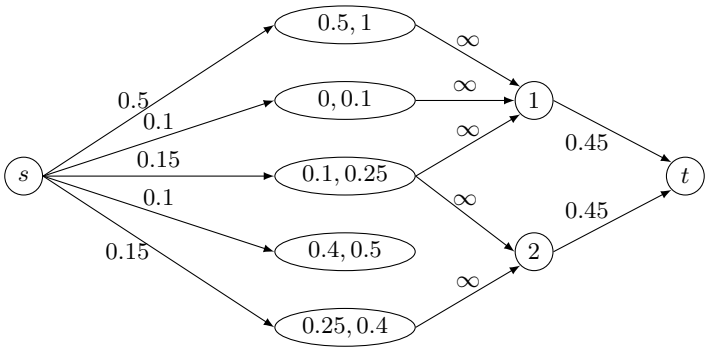

Figure 2: The flow network induced by the example.

agents; designing truthful, proportional and EF algorithms even for this case is nontrivial. Assume that $\text{len}(U_1) \leq \text{len}(U_2)$ for ease of presentation. If in addition, $\text{len}(U_1) > \text{len}(U_1 \cup U_2)/2$ then set $\{1, 2\}$ has the smallest average and we divide $U_1 \cup U_2$ exactly, with each agent getting all of $U_i \setminus U_{3-i}$ and sharing $U_1 \cap U_2$ in a way that $\text{len}(A_1) = \text{len}(A_2)$. Otherwise, agent 1 gets all of U_1 and agent 2 gets $U_2 \setminus U_1$. The algorithm tries to give both agents the same length, with each agent always getting at least half of its desired intervals, leading to proportionality and EF because of piecewise uniform valuations. For sufficient overlap in desired intervals, each receives exactly half of $U_1 \cup U_2$. For totally disjoint reference pieces, each receives just its reference piece. We defer a discussion of truthfulness to the general algorithm; the crux here is to note that each agent i receives all of $U_i \setminus U_{3-i}$, and the algorithm precludes overclaims through providing a nonincreasing share of $U_i \cap U_{3-i}$ as $\text{len}(U_i)$ increases.

Exact Allocations and Maximum Flows. Before turning to properties of truthfulness and fairness, we point out that so far it is unclear whether Algorithm 1 is well-defined. In particular, the algorithm requires an exact allocation E with respect to the subset $S_{\min}$ and X, but it remains to show that such an allocation exists, and to provide a way to compute it. To this end we exploit a close relationship between exact allocations and maximum flows in networks.

For a given set of agents $S \subseteq N$ and a piece of cake to be allocated X, define a graph $G(S, X)$ as follows. We keep track of a set of marks, which will be used to generate nodes in $G(S, X)$. First mark the left and right boundaries of all intervals that are contained in X. For each agent $i \in N$ and subinterval in U_i, mark the left and right boundaries of subintervals that are contained in $U_i \cap X$. When we have finished this process, each pair of consecutive markings will form an interval such that each agent will either uniformly value the entire interval or value none of the interval. In $G(S, X)$, create a node for each interval I formed by consecutive markings, and add a node for each agent $i \in N$, a source node s, and a sink node t. For each interval I, add a directed edge from source s to I with capacity equal to the length of the interval. Each agent node is connected to t by an edge with capacity $\text{avg}(S, X)$. For each interval-agent pair (I, i), add a directed edge with infinite capacity from node I to the agent i if agent i desires interval I.

For example, suppose $U_1 = [0, 0.25] \cup [0.5, 1]$ and $U_2 = [0.1, 0.4]$. If $X = [0, 1]$ then the interval markings will be $\{0, 0.1, 0.25, 0.4, 0.5, 1\}$. Agent 1 values $[0, 0.1]$, both agents value $[0.1, 0.25]$, agent 2 values $[0.25, 0.4]$, neither agent values $[0.4, 0.5]$ and agent 1 values $[0.5, 1]$. It holds that $\text{len}(D(\{1, 2\}, [0, 1])) = 0.9$. Average values are 0.75, 0.3 and 0.45 for sets $\{1\}, \{2\}$ and $\{1, 2\}$ respectively. See Figure 2 for an illustration of the induced flow network.

Lemma 1. *Let $S \subseteq N$, and let X be a piece of cake. There is a flow of size len$(D(S,X))$ in $G(S,X)$ if and only if for all $S' \subseteq S$, avg$(S',X) \geq$ avg(S,X).*

Below we only prove the "if" direction, which is the one we need, using an application of the classic Max-Flow Min-Cut Theorem (see, e.g., [8]).

Proof of "if". Assume that for all $S' \subseteq S$, $\text{avg}(S',X) \geq \text{avg}(S,X)$. By the Max-Flow Min-Cut Theorem, the minimum capacity removed from a graph in order to disconnect the source and sink is equal to the size of the maximum flow. The only edges with finite capacity in $G(S,X)$ are the ones that connect agent nodes to the sink, and the ones that connect the source to the interval nodes.

Construct a candidate minimum cut by disconnecting some set of agent nodes $T \subseteq S$ from the sink at cost $|T| \cdot \text{avg}(S,X)$ and then disconnecting all the (s,I) connections to interval nodes I desired by an agent $i \in S \setminus T$. This means that the total additional capacity we need to remove is $\text{len}(D(S \setminus T, X))$, the total length of intervals desired by at least one agent in $S \setminus T$. By assumption, this is at least $|S \setminus T| \cdot \text{avg}(S,X)$. As a result, this cut has capacity of at least $|T| \cdot \text{avg}(S,X) + |S \setminus T| \cdot \text{avg}(S,X) = |S| \cdot \text{avg}(S,X) = \text{len}(D(S,X))$. □

The following lemma establishes that this flow of size $\text{len}(D(S,X))$ in $G(S,X)$ is, in particular, characterizing an exact allocation. We omit the proof, which follows from the construction of the network.

Lemma 2. *Let $S \subseteq N$, and let X be a piece of cake. There exists an exact allocation with respect to S, X if and only if there exists a maximum flow of size len$(D(S,X))$ in $G(S,X)$.*

By combining Lemma 1 and Lemma 2 we see that the algorithm is indeed well-defined: if S has the smallest average then there exists an exact allocation with respect to S, X.[6] Moreover, we obtain a tractable algorithm for computing an exact allocation, by computing the maximum flow and deriving an exact allocation. A maximum flow can be computed in time that is polynomial in the number of nodes, that is, polynomial in our input size (see, e.g., [8]). We remark without proof that it is also possible to implement Step 2 of SubRoutine in polynomial time, using similar (but slightly more involved) network flow arguments. Therefore, Algorithm 1 can be implemented in polynomial time.

Truthfulness and fairness. Our main tool in proving that Algorithm 1 is truthful, proportional and EF is the following lemma (we omit its proof).

Lemma 3. *Let $S_1, \ldots, S_m$ be the ordered sequence of agent sets with the smallest average as chosen by Algorithm 1 and $X_1, \ldots, X_m$ be the ordered sequence of pieces to be allocated in calls to* SubRoutine. *That is, $X_1 = [0,1], X_2 = X_1 \setminus D(S_1, X_1), \ldots, X_m = X_{m-1} \setminus D(S_{m-1}, X_{m-1})$. Then for all $i > j$, avg$(S_i, X_i) \geq$ avg(S_j, X_j), and agents that are members of later sets receive weakly more in desired lengths.*

Envy-freeness now follows immediately from the lemma. Indeed, consider an agent $i \in N$. By "chosen" we mean that the agent was part of the subset with smallest average. The agent does not envy agents chosen in the same call to SubRoutine since all agents receive the same length in desired intervals and their valuations are piecewise uniform. By Lemma 3, the agent does not envy agents chosen in earlier calls because the amount agents receive weakly increases with each call. The agent does not envy agents chosen in later calls because all intervals desired by the agent are removed from consideration when the agent receives its allocation.

[6] Note that the network in Figure 2 does not satisfy the average minimality requirement and does not provide a corresponding exact allocation.

We provide a sketch of truthfulness, which follows by showing that an agent $i \in N$ has no incentive to change the choice of $S_{\min}$ and cannot profitably manipulate the exact allocation for a given $S_{\min}$.

1. Manipulations that change $S_{\min}$. Consider two subcases.
   (a) When i reports truthfully, $S_{\min} = S', i \notin S'$. An agent cannot affect $\text{avg}(T, X)$ if $i \notin T$, so the agent cannot cause some other $S'', i \notin S''$ to be chosen. The agent can cause $S'', i \in S''$, to be chosen, but then $\text{avg}(S'', X) \leq \text{avg}(S', X)$ and it follows from Lemma 3 that the agent does not gain.
   (b) When i reports truthfully, $S_{\min} = S', i \in S'$. Assume without loss of generality that $|S| \geq 2$. In this case, all agents in S', including i, receive exactly $\text{avg}(S', X) = k$ in intervals. Agent i can cause selection of some S'' by misstating its valuation. If $i \in S''$, then $\text{avg}(S'', X) \geq k$ for this to be profitable. If $i \notin S''$, then S'' was not chosen when i reports truthfully, so $\text{avg}(S'', X) \geq k$. In either case, agents $j \in S' \setminus \{i\}$ previously received k, but now receive at least k by observing that $\text{avg}(S'', X) \geq k$ and applying Lemma 3. Agent i receives at most $\text{len}(D(S', X))$ minus the intervals received by agents $j \in S' \setminus \{i\}$.[7] These agents receive weakly more if i manipulates, and thus, manipulations are not profitable.
2. Manipulations that change the exact allocation for a given $S_{\min}, i \in S_{\min}$. By definition each agent in $S_{\min}$ receives exactly $\text{avg}(S_{\min}, X)$ in desired intervals. If agent i decreases this value, it receives strictly less. If agent i increases this value by lying, then other agents receive more of the actual $D(S_{\min}, X)$, leaving less for agent i.

We omit the proof of proportionality, but it follows after establishing that no desired pieces are thrown away. Overall, we have the following theorem.

Theorem 4. *Assume that the agents have piecewise uniform valuation functions. Then Algorithm 1 is truthful, proportional, EF, and polynomial-time.*

Relation to work on the random assignment problem. Consider a setting where indivisible items must be assigned to agents. In the *random assignment problem* items can be assigned to agents randomly, i.e., a random assignment is a probability distribution over deterministic assignments. A random assignment that gives an item to an agent with probability p can be interpreted as assigning a p-fraction of the item to the agent. Crucially, in the papers discussed below the assumption is that each agent is only interested in receiving one item.

Bogomolnaia and Moulin [3] consider the random assignment problem when the agents have dichotomous preferences over the items, in the sense that for each agent the set of items can be partitioned into acceptable and unacceptable items (where all acceptable items have value 1 and unacceptable items have value 0). They provide a random assignment method called the *egalitarian assignment solution* and show that it is truthful, EF, and satisfies other highly desirable properties.

Interestingly, the cake cutting problem under piecewise uniform valuation functions is similar to a random assignment problem, as one can mark the beginning and end of each agent's desired intervals and treat the subintervals between consecutive marks as items. However, there are two fundamental differences between our setting and [3]. First, in our setting agents are interested in receiving as much of their desired "items" as possible (rather than just one item). Second, in our setting dichotomous preferences would mean that agents

[7] Lemma 3 also applies to agent i, but since it lies, it may receive intervals that are not desired and outside of $D(S', X)$.

value all desired subintervals equally, which is clearly not the case since these subintervals have different lengths.[8] Nevertheless, it turns out that the egalitarian assignment solution is very similar to the special case of Algorithm 1 under this strong assumption. Katta and Sethuraman [11] observe that the egalitarian assignment solution can be computed in polynomial time using network flow techniques, so our arguments above are an independent generalization of this observation. Interestingly, it is noted in [11] that the egalitarian assignment solution is identical to another independent algorithm for finding a lexicographically optimal flow in a network due to Megiddo [12].

In earlier work Bogomolnaia and Moulin [2] study random assignments under *strict* ordinal preferences, and propose a solution that satisfies a weaker notion of truthfulness (which does not imply truthfulness in our setting) as well as envy-freeness and other properties. In terms of the agents' preferences this setting is incomparable to ours since agents may be indifferent between subintervals. However, in our setting agents cannot hold arbitrary ordinal preference profiles over subintervals between consecutive marks, since if two agents desire two subintervals, both agents would value the longer subinterval more than the shorter.

The results of [2] were extended by Katta and Sethuraman [11] to the case where agents are allowed to be indifferent between items. While the assumptions of [11] regarding preferences are weaker than ours, they establish that in this more general setting even Bogomolnaia and Moulin's weaker notion of truthfulness is in fact incompatible with envy-freeness and an additional efficiency requirement; the algorithm that they propose satisfies the last two properties and hence is not (even weakly) truthful.

4 Randomized Algorithms and Piecewise Linear Valuations

In the previous section we saw that designing deterministic truthful and fair algorithms is not an easy task, even if the valuation functions of the agents are rather restricted. In this section we shall demonstrate that by allowing randomness we can obtain significantly more general results.

A *randomized cake cutting algorithm* outputs a random allocation given the reported valuation functions of the agents. There are very few previous papers regarding randomized algorithms for cake cutting. A rare example is the paper by Edmonds and Pruhs [9], where they give a randomized algorithm that achieves approximate proportionality with high probability. We are looking for a more stringent notion of fairness. We say that a randomized algorithm is *universally proportional* (resp., *universally EF*) if it always returns an allocation that is proportional (resp., EF).

One could also ask for *universal truthfulness*, that is, require than an agent may never benefit from lying, regardless of the randomness of the algorithm. A universally truthful algorithm is simply a probability distribution over deterministic truthful algorithms. However, asking for both universal fairness and universal truthfulness would not allow us to enjoy the additional flexibility that randomization provides. Therefore, we slightly relax our truthfulness requirement. Informally, we say that a randomized algorithm is *truthful in expectation* if, for all possible valuation functions of the other agents, the expected value an agent receives for its allocation cannot increase by lying, where the expectation is taken over the randomness of the algorithm.

We remark that while truthfulness in expectation seems natural, fairness (i.e., proportionality and envy-freeness) is something that we would like to hold *ex-post*; fairness is a

[8] In general no discretization of the cake would necessarily yield subintervals of equal length that correspond to dichotomous preferences. If we assume that desired intervals have rational endpoints then such a discretization can be found, but the number of subintervals would be exponentially large, leading to computational intractability.

property of the specific allocation that is being made, and continues to be relevant after the algorithm has terminated. Interestingly enough, if we were to turn this around, then achieving universal truthfulness and envy-freeness/proportionality in expectation is trivial: simply allocate the entire cake to a uniformly random agent!

4.1 A randomized algorithm

In order to design a randomized algorithm that is truthful in expectation, universally proportional, and universally EF, we consider a very special type of allocation. In the following we will not require the free disposal assumption, that is, we will consider partitions $X_1, \ldots, X_n$ of the cake such that $\bigcup_i X_i = [0, 1]$. We say that a partition $X_1, \ldots, X_n$ is *perfect* if for all $i, j \in N$, $v_i(X_j) = 1/n$. Consider the following randomized algorithm.

Algorithm 2 $(V_1, \ldots, V_n)$

1. Find a perfect partition $X_1, \ldots, X_n$.
2. Draw a random permutation π over N.
3. For each $i \in N$, set $A_i = X_{\pi(i)}$.

Lemma 5. *Algorithm 2 is truthful in expectation, universally proportional, and universally EF.*[9]

Proof. The fact that the algorithm is universally proportional and universally EF follows from the definition of perfect partitions: every agent has value $1/n$ for every piece!

We turn to truthfulness in expectation. The value an agent $i \in N$ obtains by reporting truthfully is exactly $1/n$. If agent i lies then the algorithm may choose a different partition $X'_1, \ldots, X'_n$. However, for any partition $X'_1, \ldots, X'_n$ the expected value of agent i when given a random piece is

$$\sum_{j \in N} \frac{1}{n} \cdot V_i(X'_j) = \frac{1}{n} \left(\sum_{j \in N} V_i(X'_j) \right) = \frac{1}{n},$$

where the second equality follows from the fact that the valuation functions are additive. □

Finding perfect partitions. Lemma 5 holds much promise, in that it is valid for all valuation functions. But there still remains the obstacle of actually finding a perfect partition given the valuation functions of the agents. Does such a partition exist, and can it be computed? More than two decades ago, Noga Alon [1] proved that if the valuation functions of the agents are defined by the integral of a continuous probability measure then there *exists* a perfect partition; this is a generalization of his famous theorem on necklace splitting. Unfortunately, Alon's elegant proof is nonconstructive (which is unusual for a proof in combinatorics), and to this day there is no known constructive method under general assumptions on the valuation functions. This is not surprising since a perfect partition induces an EF allocation, and finding an EF allocation in a bounded number of steps for more than four agents is an open problem.

To obtain a computational method, we consider valuation functions that are *piecewise linear*. A valuation function V_i is considered piecewise linear if and only if its corresponding value density function v_i is piecewise linear on $[0, 1]$. Piecewise linear valuation functions are significantly more general than the class of piecewise constant valuation functions. A piecewise linear valuation function can be concisely represented by the intervals on which v_i is linear, and for each interval the two parameters of the linear function. The following

[9] Mossel and Tamuz [13] make the same observation.

lemma provides us with a tractable method of finding a perfect partition when the agents have piecewise linear valuation functions.

Lemma 6. *Assume that the agents have piecewise linear valuation functions. Consider the following procedure. We make a mark at 0 and 1, and for each agent $i \in N$ make a mark at the left and right boundaries of each interval where v_i is linear. Next, we divide each interval I_j between two consecutive marks into $2n$ consecutive and connected subintervals $I_j^1, \ldots, I_j^{2n}$ of equal length. For each such I_j and every $i \in N$ add the subintervals I_j^i and I_j^{2n-i+1} to X_i. Then the overall partition is perfect.*

The lemma's proof is omitted. By combining Lemma 6 with Lemma 5 we obtain the following result.

Theorem 7. *Assume that the agents have piecewise linear valuation functions. Then there exists a randomized algorithm that is truthful in expectation, universally proportional, universally EF, and polynomial-time.*

5 Discussion

We have made progress on truthful and fair algorithms for cake cutting. In unpublished work, we show the nonexistence of simpler methods that make only contiguous allocations (and look closer to generalizations of the classic cut-and-choose algorithm) even for two agents both of whom are uniformly interested in a single (but different) subinterval. In future work we would like to generalize the deterministic algorithm to piecewise constant valuations and drop the free-disposal assumption. For practical settings, allowing more expressiveness (e.g., piecewise linear but a requirement that intervals are above some threshold length) seems important.

Acknowledgments

We thank Noga Alon for his wise advice regarding the tractable implementation of Step 2 of SubRoutine. We also thank Anna Bogomolnaia and Hervé Moulin for pointing our attention to the connection between our problem and the random assignment problem.

References

[1] N. Alon. Splitting necklaces. *Advances in Mathematics*, 63:241–253, 1987.

[2] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100:295–328, 2001.

[3] A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72:257–279, 2004.

[4] S. J. Brams, M. A. Jones, and C. Klamler. Better ways to cut a cake. *Notices of the AMS*, 53(11):1314–1321, 2006.

[5] S. J. Brams, M. A. Jones, and C. Klamler. Proportional pie-cutting. *Int. Journal of Game Theory*, 36(3–4):353–367, 2008.

[6] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution.* Cambridge University Press, 1996.

[7] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.

[8] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms.* MIT Press, 2nd edition, 2001.

[9] J. Edmonds and K. Pruhs. Balanced allocations of cake. In *Proc. of 47th FOCS*, pages 623–634, 2006.

[10] J. Edmonds and K. Pruhs. Cake cutting really is not a piece of cake. In *Proc. of 17th SODA*, pages 271–278, 2006.

[11] A. Katta and J. Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory*, 131:231–250, 2006.

[12] N. Megiddo. A good algorithm for lexicographically optimal flows in multi-terminal networks. *Bulletin of the American Mathematical Society*, 83:407–409, 1979.

[13] E. Mossel and O. Tamuz. Truthful fair division. Manuscript, 2010.

[14] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1–2):166–196, 2001.

[15] A. D. Procaccia. Thou shalt covet thy neighbor's cake. In *Proc. of 21st IJCAI*, pages 239–244, 2009.

[16] J. M. Robertson and W. A. Webb. *Cake Cutting Algorithms: Be Fair If You Can.* A. K. Peters, 1998.

[17] W. Thomson. Children crying at birthday parties. Why? *Journal of Economic Theory*, 31:501–521, 2007.

Yiling Chen
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
Email: `yiling@eecs.harvard.edu`

John K. Lai
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
Email: `jklai@post.harvard.edu`

David C. Parkes
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
Email: `parkes@eecs.harvard.edu`

Ariel D. Procaccia
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
Email: `arielpro@seas.harvard.edu`