

Simultaneous Cake Cutting

Eric Balkanski

Carnegie Mellon University, USA
ebalkans@andrew.cmu.edu

David Kurokawa

Carnegie Mellon University, USA
dkurokaw@cs.cmu.edu

Simina Brânzei

Aarhus University, Denmark
simina@cs.au.dk

Ariel D. Procaccia

Carnegie Mellon University, USA
arielpro@cs.cmu.edu

Abstract

We introduce the *simultaneous model* for cake cutting (the fair allocation of a divisible good), in which agents simultaneously send messages containing a sketch of their preferences over the cake. We show that this model enables the computation of divisions that satisfy *proportionality* — a popular fairness notion — using a protocol that circumvents a standard lower bound via parallel information elicitation. Cake divisions satisfying another prominent fairness notion, envy-freeness, are impossible to compute in the simultaneous model, but admit arbitrarily good approximations.

Introduction

The theory of fair division provides formal notions of fairness, and mechanisms for computing outcomes that achieve these notions. When the good to be shared is divisible and heterogeneous, the problem of fairly dividing it is widely referred to as *cake cutting*. But the childish name is misleading, as this is a problem of great mathematical depth and serious potential applications, which, for many decades, has attracted thinkers from mathematics, economics, and political science; see, e.g., the books by Brams and Taylor (1996), Robertson and Webb (1998), and the survey by Procaccia (2013).

Cake cutting is largely an algorithmic endeavor, but only recently have computer scientists started to weigh in. One of the exciting developments is the emergence of a computational model (formally, a concrete complexity model) of cake cutting that is attributed to Robertson and Webb (1998). This model views the cake cutting process as an interaction between the protocol and the agents, where the protocol elicits information about the agents' valuations by iteratively asking two types of queries: either cutting a piece of certain value, or evaluating a given piece. The Robertson-Webb model allows computer scientists to formally reason about the complexity of achieving fair cake divisions (Edmonds and Pruhs 2006b; 2006a; Woeginger and Sgall 2007; Procaccia 2009; Deng, Qi, and Saberi 2012; Aumann, Dombb, and Hassidim 2013; Kurokawa, Lai, and Procaccia 2013; Brânzei, Procaccia, and Zhang 2013).

In this paper, we introduce a novel computational model that, we believe, provides a fundamentally new perspective

on cake cutting; we call it the *simultaneous model*. In our model, the agents *simultaneously* report compact versions of their preferences, namely, their values for specific pieces of cake; this information is used to compute a fair allocation, without further communication among the agents. We define the *complexity* of a simultaneous protocol as the maximum number of pieces whose values an agent may need to report.¹

Cake cutting protocols in the simultaneous model have two advantages compared to their counterparts in the Robertson-Webb model:

1. Elicitation of preferences can be done in parallel.
2. The existence of computationally efficient simultaneous protocols would imply that agents' valuation functions can be *sketched* in a way that preserves sufficient information for recovering a fair cake division (via the protocol).

On the other hand, the simultaneous model severely restricts the power of protocols. Is the restriction so severe that fair divisions, according to standard fairness properties, cannot be computed? Our research question is

... which fairness properties are computationally feasible in the simultaneous model, and what is the complexity of computing cake divisions satisfying those properties?

Our Results

Our results focus on the two most widely-studied fairness properties, *proportionality* and *envy-freeness*.

Proportionality requires that each agent receives a piece of value at least $1/n$ of the entire cake (according to that agent's preferences), where n is the number of agents. We show that the complexity of proportionality in the simultaneous model is exactly n , i.e., there exists a proportional simultaneous protocol with complexity n (each agent reports its value for at most n pieces of cake), and no proportional simultaneous protocol can do better. In contrast, it is known that the complexity of proportional cake cutting in the Robertson-Webb model is $\Theta(n \log n)$ (Edmonds and Pruhs 2006b). The

¹This definition is better, formally and intuitively, than taking the *overall* amount of communication (summed over all agents); it is also consistent with related work on communication complexity (Kremer, Nisan, and Ron 1999).

reduction in complexity below the formal Robertson-Webb lower bound is driven by parallel information elicitation.

Envy-freeness stipulates that each agent (weakly) prefers his allocation to the allocation of any other agent. While envy-free allocations can always be computed in the Robertson-Webb model (Brams and Taylor 1995), we show there exists no simultaneous protocol that can guarantee exact envy-free allocations for every instance of the valuations. However, we devise a simultaneous protocol that guarantees allocations arbitrarily close to envy-freeness.

Related Work

In recent years, cake cutting has emerged as a major research topic in artificial intelligence (Procaccia 2009; Caragiannis, Lai, and Procaccia 2011; Cohler et al. 2011; Brams et al. 2012; Bei et al. 2012; Aumann, Dombb, and Hassidim 2013; Kurokawa, Lai, and Procaccia 2013; Brânzei, Procaccia, and Zhang 2013; Brânzei and Miltersen 2013; Chen et al. 2013). The growing interest in cake cutting, and fair division more broadly, is partly motivated by potential applications in AI, such as industrial procurement, manufacturing and scheduling, and airport traffic management (Chevalyre et al. 2006). For example, concrete applications to the allocation of multiple *computational* resources in shared computing systems have recently received significant attention (Gutman and Nisan 2012; Kash, Procaccia, and Shah 2013). The growing interest is also driven by a newfound understanding that computational thinking can provide a completely new perspective on this classical research topic.

Our simultaneous model of cake cutting is related to, and conceptually draws on, work on *communication complexity* (Kushilevitz and Nisan 1996) and *streaming algorithms* (Muthukrishnan 2005). In particular, Kremer et al. (1999) studied the relation between one-round communication complexity and simultaneous communication complexity. Similarly to our model, streaming algorithms deal with compact representations — called *sketches* — of data. Some papers focus specifically on sketching valuation functions or preferences in various contexts (Bachrach, Porat, and Rosenschein 2009; Caragiannis and Procaccia 2011; Badanidiyuru et al. 2012).

Cake Cutting Background

The cake is modeled as the interval $[0, 1]$; there is also a set $N = \{1, \dots, n\}$ of agents. A *piece of cake* X is a finite set of disjoint intervals of $[0, 1]$.

Each agent is endowed with an integrable, non-negative value density function $v_i(x)$ that induces a value for each possible piece of cake. Formally, the value of agent i for a piece X is given by:

$$V_i(X) = \sum_{I \in X} \int_I v_i(x) dx.$$

By definition, the valuations of the agents are additive, i.e. $V_i(X \cup Y) = V_i(X) + V_i(Y)$ if X and Y are disjoint; and non-atomic, i.e. $V_i([x, x]) = 0$. We assume that each agent has a value of one for the entire cake: $V_i([0, 1]) = 1$ for all $i \in N$. This assumption is without loss of generality for the purposes of this paper.

An *allocation* $A = (A_1, \dots, A_n)$ is a partition of the cake among the agents, that is, each agent i receives the piece A_i , the pieces are disjoint, and $\bigcup_{i \in N} A_i = [0, 1]$.

An allocation A is *Pareto optimal* if there does not exist another allocation A' such that $V_i(A'_i) \geq V_i(A_i)$ for all $i \in N$, and the inequality is strict for at least one agent. Pareto optimality is a notion of economic efficiency.

Two other important properties concern fairness. An allocation A is *proportional* if for all $i \in N$, $V_i(A_i) \geq 1/n$, and *envy-free* if for all $i, j \in N$, $V_i(A_i) \geq V_i(A_j)$. Envy-freeness implies proportionality when the entire cake is allocated, but the converse is not always true.

The standard model of communication between the center and agents in cake cutting was proposed by Robertson and Webb (1998), and employed in a body of work studying the complexity of cake cutting (Edmonds and Pruhs 2006b; 2006a; Woeginger and Sgall 2007; Procaccia 2009; Kurokawa, Lai, and Procaccia 2013). The model restricts the interaction between the protocol and the agents to two types of queries:

- *Cut* query: $\text{Cut}_i(x, \alpha)$ asks agent i to return a point y such that $V_i([x, y]) = \alpha$.
- *Evaluate* query: $\text{Eval}_i(x, y)$ asks agent i to return a value α such that $V_i([x, y]) = \alpha$.

We illustrate the Robertson-Webb model using the most basic cake cutting protocol, *Cut and Choose*, which computes an envy-free allocation for two agents. Under this protocol, agent 1 cuts the cake into two pieces that it values equally, and agent 2 selects its preferred piece, leaving the remaining piece for agent 1. The *Cut and Choose* protocol can be simulated in the Robertson-Webb model using two queries: $y = \text{cut}_1(0, \frac{1}{2})$, and $\text{eval}_2(0, y)$. The first query creates two pieces $X = [0, y]$ and $Y = [y, 1]$ such that $V_1(X) = V_1(Y) = 1/2$. The second query gives $V_2(X)$. At this point, there is enough *information* to pinpoint an envy-free allocation: If the answer to the second query is at least $1/2$, we can allocate X to agent 2 and Y to agent 1, otherwise we swap the pieces. We emphasize that a protocol in the Robertson-Webb model does not actually return an allocation; the goal is to elicit enough information such that there exists an allocation that is guaranteed to satisfy a given property (such as envy-freeness).

The Simultaneous Model

We define a *discretization* of the cake as a tuple (\bar{x}, \bar{w}) , for which there exists $m \in \mathbb{N}$ such that:

- $\bar{x} = (x^0, x^1, \dots, x^{m-1}, x^m)$ is a sequence of *cut points* with $0 = x^0 < x^1 < \dots < x^{m-1} < x^m = 1$.
- $\bar{w} = (w^1, \dots, w^m)$ is a sequence of values, such that w^i represents the value of the piece $[x^{i-1}, x^i]$ and $w^1 + \dots + w^m = 1$.

Let \mathcal{D} denote the space of all discretizations. Then a one-round protocol can be defined as follows:

Definition 1 (Simultaneous protocol). A *simultaneous protocol* is a function $\mathcal{F} : \mathcal{V} \rightarrow \mathcal{D}$, where \mathcal{V} is the space of valuations, \mathcal{D} is the space of discretizations of the cake, and

$\mathcal{F}(V)$ is the discretization that an agent is instructed to report when his valuation function is V .

One could alternatively define a simultaneous protocol as reporting a set of (possibly overlapping) subintervals and their values. However, the two definitions are essentially equivalent for our purposes.

What does it mean for a simultaneous protocol to satisfy a *property*, such as envy-freeness, proportionality, or Pareto optimality? This question involves surprising subtleties even in the Robertson-Webb model, and so the definition must be carefully chosen. Very roughly speaking, the main difficulty (in both models) is that agents could potentially use an injection from the space of valuation functions to $[0, 1]$ to encode their entire valuation function as a single number (e.g., the first cut point they make). For any given property, that would give enough information to compute an allocation satisfying the property (if one exists). The definition below circumvents this problem, by capturing the idea that reporting a value for an interval commits the agent to a valuation function that actually assigns the reported value to that interval, and nothing else.

Definition 2 (Property of a simultaneous protocol). Let \mathcal{P} be a property of cake allocations. A protocol \mathcal{F} satisfies property \mathcal{P} if the following holds for any tuple of valuation functions $\bar{V} = (V_1, \dots, V_n)$:

- Whenever each agent i follows the protocol by reporting its recommended discretization, $(\bar{x}_i, \bar{w}_i) := \mathcal{F}(V_i)$, there exists an allocation A that satisfies \mathcal{P} with respect to any other valuations $\bar{V}' = (V'_1, \dots, V'_n)$ consistent with the discretizations reported at \bar{V} (i.e., $V'_i([x_i^{j-1}, x_i^j]) = w_i^j, \forall i, j$).

For example, let us describe an envy-free simultaneous protocol \mathcal{F} for two agents. $\mathcal{F}(V_i)$ is the discretization (\bar{x}_i, \bar{w}_i) where $\bar{x}_i = (x_i^0 = 0, x_i^1, x_i^2 = 1)$ and $\bar{w}_i = (\frac{1}{2}, \frac{1}{2})$, that is, each agent essentially cuts the cake into two pieces worth $1/2$ using the cut point x_i^1 . Now assume without loss of generality that $x_1^1 \leq x_2^1$, and consider the allocation $A_1 = [0, x_1^1]$, $A_2 = [x_1^1, 1]$. This allocation is clearly envy-free for the reported valuation functions, and, moreover, it is envy free for any valuation functions where $V'_i([0, x_i^1]) = V'_i([x_i^1, 1]) = 1/2$ for $i = 1, 2$.

Note that in the classical setting, properties such as envy-freeness and proportionality are always satisfied from the point of view of an agent that follows the protocol sincerely, regardless of the strategies of other agents. In contrast, Definition 2 seems weaker, as it requires that all agents report valid discretizations in order for a property to be satisfied. However, this definition is compatible with properties that must be defined for all agents simultaneously, such as Pareto optimality.

In the Robertson-Webb model, the complexity of a protocol is the maximum number of cut and evaluation queries. We use an equivalent definition in the simultaneous model.

Definition 3 (Complexity of a simultaneous protocol). The *complexity* of a simultaneous protocol is the maximum number of intervals in the discretization $\mathcal{F}(V)$ taken over all $V \in \mathcal{V}$ (that is, the maximum number of cut points minus

one). If the maximum does not exist, we say that the protocol is *unbounded*.

For example, the complexity of the envy-free simultaneous protocol for two agents is 2.

Proportionality

We start by examining proportionality in the simultaneous model. In the Robertson-Webb query model, the complexity of computing proportional allocations is $\Theta(n \log n)$: an $O(n \log n)$ upper bound is given by the Even-Paz (1984) protocol, and a matching lower bound was established by Edmonds and Pruhs (2006b).

Similarly, the simultaneous model turns out to admit the computation of proportional allocations, but the complexity of proportionality in this model is only $\Theta(n)$. For the upper bound, we describe a protocol that is a simultaneous interpretation of a protocol designed in a different context by Manabe and Okamoto (2012). Importantly, this protocol requires $\Theta(n^2)$ cut queries in the Robertson-Webb model; but the simultaneous model allows us to implicitly parallelize the queries to the agents, leading to a reduction in complexity. The simultaneous model captures the insight that the information elicited from one agent does not need to rely on the information elicited from another.

Theorem 1. *There exists a proportional simultaneous protocol with complexity n .*

Proof. Consider the following simultaneous protocol:

- *Map the valuation function of each agent to n disjoint contiguous intervals of value exactly $1/n$ each.*

Formally, the discretization is defined by $\bar{x}_i = (x_i^0, \dots, x_i^n)$ and $w_i^j = 1/n$, for all $j = 1, \dots, n$.

Given the intervals submitted by the agents, we produce an allocation by scanning the cake from left to right until the first mark, $x_{i_1}^1$, of some agent $i_1 \in N$ is encountered. Allocate the piece $[0, x_{i_1}^1]$ to agent i_1 . Then, scan to the right starting with the point $x_{i_1}^1$ while looking for the second mark $x_{i_2}^2$ of some agent $i_2 \in N \setminus \{i_1\}$. Allocate the piece $[x_{i_1}^1, x_{i_2}^2]$ to agent i_2 and continue in this fashion until the entire cake is allocated.

To see why the protocol is proportional, note that for agent i_t that was allocated in round t , $x_{i_t}^{t-1} \geq x_{i_{t-1}}^{t-1}$, because i_t was not selected in round $t-1$. Thus, $[x_{i_t}^{t-1}, x_{i_t}^t] \subseteq [x_{i_{t-1}}^{t-1}, x_{i_t}^t]$. Moreover, $A_{i_t} = [x_{i_{t-1}}^{t-1}, x_{i_t}^t]$ and $V_{i_t}([x_{i_t}^{t-1}, x_{i_t}^t]) = 1/n$, thus $V_{i_t}(A_{i_t}) \geq 1/n$. \square

Next, we show the bound given in Theorem 1 is tight.

Theorem 2. *Every proportional simultaneous protocol has complexity at least n .*

Proof. Assume by contradiction that there exists a proportional simultaneous protocol \mathcal{F} with complexity less than n . Without loss of generality, let V_n be a valuation function such that $\mathcal{F}(V_n)$ reports the values of $n-1$ intervals with cut points $(x_n^0, \dots, x_n^{n-1})$. (The case where the

agent reports fewer intervals is similar.) Then the valuations of the other agents can be set such that for every agent $i \in N \setminus \{n\}$, the entire value of the cake from the point of view of agent i is concentrated in the interval $[x_n^{i-1}, x_n^i]$, that is, $V_i([x_n^{i-1}, x_n^i]) = 1$.

Let us now consider two (exhaustive) types of allocations. First, let A be an allocation such that for all $i \in \{1, \dots, n-1\}$, agent i gets a nonempty interval $I_i \subseteq [x_n^{i-1}, x_n^i]$. We can define the valuation function V'_n where $V'_n(I_i) = V_n([x_n^{i-1}, x_n^i])$ for all $i \in \{1, \dots, n-1\}$. Then V'_n is consistent with agent n 's reported intervals, but $V'_n(A_n) = 0$, so the allocation is not proportional with respect to V'_n .

Second, let A be an allocation such that there exists an agent $i \in \{1, \dots, n-1\}$ that does not get a nonempty interval $I_i \subseteq [x_n^{i-1}, x_n^i]$. Then clearly $V_i(A_i) = 0$, and again the allocation is not proportional. \square

(Approximate) Envy-Freeness, and Beyond

We have seen that simultaneous protocols can compute proportional allocations. For two agents, proportionality and envy-freeness coincide, but for more agents, envy-freeness is strictly stronger. It has long been known that envy-free allocations are guaranteed to exist, but it wasn't until the nineties that an envy-free protocol that can be simulated in the Robertson-Webb model was discovered (Brams and Taylor 1995).

The Brams-Taylor protocol is finite (i.e. terminates on every instance), but unbounded: its running time cannot be bounded by a function of the number of agents, and so the execution can take arbitrarily long depending on the valuation functions themselves. It is an open problem whether a bounded envy-free protocol exists in the Robertson-Webb model for any number of agents.

Next we show there exists no simultaneous protocol that is envy-free. Interestingly, this impossibility result does not assume the protocol is bounded: it says there are valuation functions for which there is no fine enough discretization to guarantee envy-freeness in the simultaneous model.

Theorem 3. *For $n \geq 3$ there does not exist an envy-free simultaneous protocol.*

Proof. Let V_1 be the uniform valuation function (i.e., its value density function is $v(x) \equiv 1$), which yields a discretization $\mathcal{F}(V_1) = (\bar{x}_1, \bar{w}_1)$ under protocol \mathcal{F} . Let there be m reported intervals, and denote $X^i = [x_1^{i-1}, x_1^i]$ for $i = 1, \dots, m$; then $w_1^i = |X^i| = x_1^i - x_1^{i-1}$. We will show that there exist valuation functions for the other agents such that no envy-free allocation can be computed from these reported intervals.

Define a constant $c \in \left(\frac{1}{w_1^1+1}, 1\right)$ such that for all $i \in N \setminus \{1\}$, the value density function v_i of agent i satisfies the following conditions:

- (a) For all $j \in \{1, \dots, m\}$, v_i is constant on X^j .
- (b) $V_i(X^1) = c \cdot w_1^1 + 1 - c$
- (c) $V_i(X^j) = c \cdot w_1^j$, for all $j \in \{2, \dots, m\}$

- (d) There do not exist distinct indices $a_1, \dots, a_x \in \{1, \dots, m\}$ such that the following identity holds:

$$w_1^{a_1} + \dots + w_1^{a_x} = \frac{1}{c \cdot n}.$$

Note that any $c \in \left(\frac{1}{w_1^1+1}, 1\right)$ induces valid valuation functions that satisfy (b) and (c), because

$$V_i([0, 1]) = \sum_{j=1}^m V_i(X^j) = c \left(\sum_{j=1}^m w_1^j \right) + (1 - c) = 1.$$

Moreover, constraint (d) can be satisfied because there is an (uncountably) infinite number of possible values of c , and the constraint only rules out a finite number of them.

Let $A = (A_1, \dots, A_n)$ be an allocation computed by the protocol. We consider two cases, depending on whether the interval X^1 is split or not among the agents.

Case I: Interval X^1 is not split. We have several subcases:

- (i) $|A_1| < \frac{1}{n}$: Then there exists another agent i that receives a piece of length at least $\frac{1}{n}$ and agent 1 envies i .
- (ii) $|A_1| \geq \frac{1}{n}$ and agent 1 receives X^1 . Then the value of the other agents for the piece received by 1 is:

$$\begin{aligned} c \cdot w_1^1 + 1 - c + c(|A_1| - w_1^1) &= c \cdot |A_1| + 1 - c \\ &\geq \frac{c}{n} + 1 - c. \end{aligned}$$

The length of the piece for all the other agents is at most $\frac{n-1}{n}$. Since the remainder of the cake does not contain X^1 , the minimum value $V_i(A_i)$ — taken over all $i \in \{2, \dots, n\}$ — is at most $c \left(\frac{n-1}{n}\right) \left(\frac{1}{n-1}\right) = \frac{c}{n}$. It follows that there exists an agent i that envies 1.

- (iii) $|A_1| \geq \frac{1}{n}$ and an agent $i \in N \setminus \{1\}$ receives X^1 . It must be the case that $|A_j| = |A_k|$ for all $j, k \in N \setminus \{i\}$ to prevent envy. For the same reason, we have $V_j(A_i) = V_j(A_j)$ for all $j \in \{2, \dots, n\}$. Therefore, all the agents, except agent 1, value A_1, \dots, A_n equally. It follows that $V_j(A_1) = \frac{1}{n}$ for all $j \in \{2, \dots, n\}$. This implies that $|A_1| = \frac{1}{c \cdot n}$, so by property (d), there exists a reported interval that is split between agent 1 and at least one other agent. Now we can define V'_1 that is consistent with $\mathcal{F}(V_1)$, where agent 1's value for its part(s) of the split interval(s) is zero; then agent 1 would be envious.

Case II: Interval X^1 is split among at least two agents. For each $i \in N$, let $A'_i = A_i \setminus X^1$. We have two subcases:

- (i) There exists exactly one agent $i \in N \setminus \{1\}$, such that $A_i \cap X^1 \neq \emptyset$. Then $A_1 \cap X^1 \neq \emptyset$. Consider another agent $j \in N \setminus \{1, i\}$.
 - If $|A_j| > |A'_1|$, let V'_1 be a valuation function consistent with $\mathcal{F}(V)$ such that agent 1 has a value of zero for his portion of X^1 . Then $V'_1(A_1) = V'_1(A'_1) < V'_1(A_j)$, violating envy-freeness.

- If $|A_j| \leq |A'_1|$, then $V_j(A'_1) \geq V_j(A_j)$. Moreover, $V_j(A_1 \setminus A'_1) > 0$. It follows that $V_j(A_1) > V_j(A_j)$, violating envy-freeness.

(ii) There exist distinct agents $i, j \in N \setminus \{1\}$ such that $A_i \cap X^1 \neq \emptyset$, $A_j \cap X^1 \neq \emptyset$. Assume without loss of generality that agent i 's piece satisfies $|A_i \cap X^1| \leq \frac{|X^1|}{2}$. Then $V_i(A_i \cap X^1) \leq \frac{1}{2}(c \cdot w_1^1 + 1 - c)$ and $V_i(A_1) \geq c|A'_1|$. It must also be the case that

$$|A'_i| \geq |A'_1| - \frac{1}{2} \left(w_1^1 + \frac{1}{c} - 1 \right)$$

since otherwise:

$$\begin{aligned} V_i(A_i) &\leq c|A'_i| + \frac{1}{2} (c \cdot w_1^1 + 1 - c) \\ &< c \left(|A'_1| - \frac{1}{2} \left(w_1^1 + \frac{1}{c} - 1 \right) \right) + \frac{1}{2} (c \cdot w_1^1 + 1 - c) \\ &= c|A'_1|, \end{aligned}$$

which implies $V_i(A_1) > V_i(A_i)$. Let V'_1 be consistent with $\mathcal{F}(V_1)$, such that $V'_1(A_1 \cap X^1) = 0$ and $V'_1(A_i \cap X^1) = w_1^1$. Then agent 1's value for i 's piece is:

$$\begin{aligned} V'_1(A_i) &= w_1^1 + |A'_i| \geq w_1^1 + |A'_1| - \frac{1}{2} \left(w_1^1 + \frac{1}{c} - 1 \right) \\ &> w_1^1 + |A'_1| - \frac{1}{2} (w_1^1 + w_1^1) \\ &= |A'_1| = V'_1(A_1), \end{aligned}$$

where the third transition holds by the choice of $c \in \left(\frac{1}{w_1^1 + 1}, 1 \right)$. Thus agent 1 envies agent i . \square

Theorem 3 tells us that we cannot hope to obtain envy-free allocations in the simultaneous model. However, it turns out that we can reach envy-free allocations arbitrarily close. Indeed, we say that an allocation is ϵ -envy free if for all $i, j \in N$, $V_i(A_i) \geq V_i(A_j) - \epsilon$. This notion of approximate envy-freeness has been studied in several previous papers (Lipton et al. 2004; Cohler et al. 2011; Deng, Qi, and Saberi 2012). We will show there exists an ϵ -envy-free protocol of polynomial complexity in n and ϵ .

The main idea is to sketch the agents' valuations using a very fine discretization, but then use a coarser discretization to partition the cake into *indivisible* goods. Then, each agent's value for each indivisible good can be accurately estimated using the fine discretization. An allocation of the indivisible goods that is approximately envy-free with respect to the estimated values is therefore also approximately envy-free with respect to the real values (with a slightly worse additive approximation term).

Theorem 4. *For every $\epsilon > 0$ there exists an ϵ -envy-free simultaneous protocol with complexity $O(n/\epsilon^2)$.*

The proof uses the following lemma, which is a special case of a result by Lipton et al. (2004), and deals with the allocation of *indivisible* goods. In this context, the valuation functions are said to be *additive* if the value of a bundle of goods is the sum of values of goods in the bundle.

Lemma 1 (Lipton et al. 2004). *Let V'_1, \dots, V'_n be additive valuation functions over a set G of indivisible goods. Assume that for all $i \in N$ and $g \in G$, $V_i(g) \leq \epsilon$. Then there exists an ϵ -envy-free allocation.*

Proof of Theorem 4. For every n and $\epsilon > 0$ we design a simultaneous protocol $\mathcal{F}^{n, \epsilon}$. Given a valuation V , $\mathcal{F}^{n, \epsilon}$ discretizes the cake as follows. First, the *coarse discretization* has $1/\delta$ subintervals of value δ each, for $1/\delta = \lceil 2/\epsilon \rceil$; note that $\delta \leq \epsilon/2$. Second, the *fine discretization* includes $1/\delta'$ intervals of value δ' each, for $1/\delta' = \lceil 16n/\epsilon^2 \rceil$; note that $\delta' \leq \epsilon^2/16n$. Formally speaking, $\mathcal{F}^{n, \epsilon}(V)$ contains the cut points of both discretizations, but we prefer to think of these two different discretizations for ease of exposition.

Given $\mathcal{F}^{n, \epsilon}(V_1), \dots, \mathcal{F}^{n, \epsilon}(V_n)$, we wish to show there is an allocation A that is ϵ -envy free with respect to any valuation functions that are consistent with these reported discretizations. Consider the partition of the cake obtained by ordering the cut points of all agents' coarse discretizations, and treating the subinterval between two adjacent cut points as an *indivisible good*. Denote the set of indivisible goods by G . This partition into indivisible goods has two properties:

1. For each indivisible good $g \in G$, $V_i(g) \leq \epsilon/2$ for all $i \in N$, because for each $i \in N$ there is a subinterval of the coarse discretization of V_i that contains g .
2. The number of indivisible goods is given by the number of "internal" (not 0 or 1) cut points plus one, i.e.,

$$|G| \leq n \left(\left\lceil \frac{1}{\delta} \right\rceil - 1 \right) + 1 \leq \frac{4n}{\epsilon}.$$

Let us create additive valuation functions V'_1, \dots, V'_n over the indivisible goods in G . For $g \in G$, let $H_i(g)$ be the set of intervals in the fine partition of V_i that are contained inside g . Define $V'_i(g) = \delta' \cdot |H_i(g)|$. We claim that

$$V'_i(g) \leq V_i(g) \leq V'_i(g) + 2\delta'. \quad (1)$$

Indeed, the left inequality is trivial. For the right inequality, let I be the interval obtained by taking $H_i(g)$ and adding one subinterval to the left and one to the right. It holds that $g \subseteq I$, hence

$$V_i(g) \leq V_i(H_i(g)) = \delta' \cdot (|H_i(g)| + 2) = V'_i(g) + 2\delta'.$$

Note that for all $i \in N$ and $g \in G$, $V'_i(g) \leq V_i(g) \leq \epsilon/2$. We can therefore use Lemma 1 to create an allocation A of the goods G such that for all $i, j \in N$, $V'_i(A_i) \geq V'_i(A_j) - \epsilon/2$. We claim that A is ϵ -envy free with respect to the original valuation functions (and any other valuations that are consistent with the reported discretizations). Indeed,

$$\begin{aligned} V_i(A_i) &\geq V'_i(A_i) \geq V'_i(A_j) - \frac{\epsilon}{2} \\ &= \left(\sum_{g \in A_j} V'_i(g) \right) - \frac{\epsilon}{2} \geq \left(\sum_{g \in A_j} (V_i(g) - 2\delta') \right) - \frac{\epsilon}{2} \\ &= V_i(A_j) - 2\delta'|A_j| - \frac{\epsilon}{2} \geq V_i(A_j) - 2\delta'|G| - \frac{\epsilon}{2} \\ &\geq V_i(A_j) - 2 \cdot \frac{\epsilon^2}{16n} \cdot \frac{4n}{\epsilon} - \frac{\epsilon}{2} = V_i(A_j) - \epsilon. \end{aligned}$$

where the first and fourth transitions follow from Equation (1). \square

Envy-freeness and proportionality are examples of what we call *linear* properties, in the sense that they are specified by linear constraints involving the agents' valuations for pieces. Another example of a linear property is *equitability*, which requires that $V_i(A_i) = V_j(A_j)$ for all $i, j \in N$, that is, agents must have identical values for their own pieces. We formally define linear properties using the matrix form, as is done in linear programs.

Definition 4 (Linear property). A property of allocations is *linear* if there exist $m \in \mathbb{N}$, matrix $B \in \mathbb{R}^{m \times n^2}$ such that $\sum_{j=1}^{n^2} |B_{ij}| \leq 1$ for $i = 1, \dots, m$, and vector $c \in \mathbb{R}^m$, such that an allocation A satisfies the property if it satisfies the constraints: $B \cdot \alpha \geq c$, where $\alpha_k = V_i(A_j)$, with $i = \lfloor \frac{k}{n} \rfloor$, and $j = k \bmod n$ if $n \nmid k$ and $j = n$ otherwise.

To illustrate the definition of α_k , note that $\alpha_1 = V_1(A_1)$, $\alpha_n = V_1(A_n)$, and $\alpha_{n+1} = V_2(A_1)$. Importantly, this representation captures equality constraints, as they can be represented using two inequalities. Furthermore, the assumption that $\sum_{j=1}^{n^2} |B_{i,j}| \leq 1$ is without loss of generality: we just divide each entry in the matrix B and vector c by the maximum sum of absolute values of any row of B , which is a constant in the context of the properties we are interested in.

For example, consider the representation of envy-freeness as a linear property for 3 agents. Let $m := n(n-1) = 6$ and define:

$$B = \frac{1}{2} \cdot \begin{pmatrix} 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 \end{pmatrix}$$

and

$$c = (0, 0, 0, 0, 0, 0).$$

Then the constraint $B \cdot \alpha \geq c$ is equivalent to requiring that $V_i(A_i) - V_i(A_j) \geq 0$, for each $i \neq j$.

Every linear property P (defined by a matrix B and a vector c) naturally admits an approximate version \mathcal{P}_ϵ , which requires each linear constraint of \mathcal{P} to hold up to an error of ϵ ; formally, $B \cdot \alpha \geq c - \epsilon \cdot \vec{1}$. Using this new notion, we can establish a more general version of Theorem 4.

Theorem 5. *For every $\epsilon > 0$ and every bounded protocol in the Robertson-Webb model that allocates the entire cake and guarantees some linear property \mathcal{P} with complexity $f(n)$, there exists a simultaneous protocol that guarantees the property \mathcal{P}_ϵ with complexity $O(f(n)/\epsilon)$.*

The proof appears in the full version of the paper (it also contains a formal definition of properties in the Robertson-Webb model). Theorem 5 implies Theorem 4 because ϵ -envy free allocations can be computed in the Robertson-Webb model using $O(n/\epsilon)$ queries. And while exact equitability is impossible to achieve in the Robertson-Webb model (Chechlárová and Pillárová 2012), ϵ -equitability can also be achieved with $O(n/\epsilon)$ queries, leading to an ϵ -equitable simultaneous protocol with complexity $O(n/\epsilon^2)$.

A technique for approximating general density functions with piecewise constant density functions (Cohler et al. 2011, Lemma 8) can be leveraged to obtain a strictly weaker version of Theorems 4 and 5, requiring the assumption that the value density functions are piecewise K -Lipschitz continuous, and giving a bound that also depends on K .

Discussion

In some ways, simultaneous protocols are weaker than their counterparts in the Robertson-Webb model: agents cannot interact, but rather are allowed to send one message only. However, in other ways simultaneous protocols are stronger. Indeed, under the Robertson-Webb model, information is elicited via cut and evaluation queries, without ever seeing the full valuations. This means that properties such as Pareto optimality are impossible to achieve in this model, even when the value density functions are restricted to be piecewise constant and the protocol is allowed to have unbounded complexity (Kurokawa, Lai, and Procaccia 2013). Intuitively, the reason is that a Pareto optimal allocation cannot allocate to agent i a subinterval I such that $V_i(I) = 0$ and $V_j(I) > 0$. But in the Robertson-Webb model, it is impossible to exactly identify the boundaries of subintervals that are worthless to an agent.

In contrast, in the simultaneous model agents can observe their full valuations before deciding the subintervals to report, which allows them to exactly mark worthless intervals. Suppose for simplicity that the agents' value density functions are piecewise constant, so each has a finite number of intervals on which its density is zero. Each agent reports a discretization that pinpoints the zero-density intervals. Then we can allocate the intervals using *serial dictatorship*: in stage i , allocate to agent i all unclaimed intervals on which its density is positive. This allocation is clearly Pareto-optimal.

Unfortunately, the protocol just described is not formally Pareto optimal according to Definition 2, because the allocation is not guaranteed to be Pareto optimal with respect to all valuation functions consistent with the reports (some may have additional worthless subintervals). In fact, Pareto optimality cannot be guaranteed in the simultaneous model — as can be shown using an argument that is similar to the proof of the equivalent result in Robertson-Webb (Kurokawa, Lai, and Procaccia 2013, Theorem 5). However, this negative result can be circumvented via a slight augmentation of the simultaneous model, which allows agents to mark intervals on which their density is strictly positive.

It is therefore natural to consider a relaxed model that allows protocols to enjoy the best of both worlds: multi-round interaction à la Robertson-Webb, and allowing agents to report discretizations by observing their own valuation function (and information previously communicated by others) à la the simultaneous model. This hybrid model gives rise to intriguing questions. Most importantly: does it admit *bounded* envy-free protocols? We view this question as a natural, compelling relaxation of what is perhaps the most enigmatic open problem in computational fair division (Procaccia 2013): settling the existence of bounded envy-free protocols in the Robertson-Webb model.

Acknowledgements

Simina Brânzei was supported by the Sino-Danish Center for the Theory of Interactive Computation, funded by the Danish National Research Foundation and the National Science Foundation of China (under the grant 61061130540), and by the Center for research in the Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council. Simina also acknowledges support from an IBM Ph.D. fellowship. David Kurokawa and Ariel D. Procaccia were supported in part by the NSF grant CCF-1215883.

References

- Aumann, Y.; Dombb, Y.; and Hassidim, A. 2013. Computing socially-efficient cake divisions. In *Proc. of 12th AAMAS*, 343–350.
- Bachrach, Y.; Porat, E.; and Rosenschein, J. S. 2009. Sketching techniques for collaborative filtering. In *Proc. of 21st IJCAI*, 2016–2021.
- Badanidiyuru, A.; Dobzinski, S.; Fu, H.; Kleinberg, R.; Nisan, N.; and Roughgarden, T. 2012. Sketching valuation functions. In *Proc. of 23rd SODA*, 1025–1035.
- Bei, X.; Chen, N.; Hua, X.; Tao, B.; and Yang, E. 2012. Optimal proportional cake cutting with connected pieces. In *Proc. of 26th AAI*, 1263–1269.
- Brams, S. J., and Taylor, A. D. 1995. An envy-free cake division protocol. *The American Mathematical Monthly* 102(1):9–18.
- Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Brams, S. J.; Feldman, M.; Morgenstern, J.; Lai, J. K.; and Procaccia, A. D. 2012. On maxsum fair cake divisions. In *Proc. of 26th AAI*, 1285–1291.
- Brânzei, S., and Miltersen, P. B. 2013. Equilibrium analysis in cake cutting. In *Proc. of 12th AAMAS*, 327–334.
- Brânzei, S.; Procaccia, A. D.; and Zhang, J. 2013. Externalities in cake cutting. In *Proc. of 23rd IJCAI*, 55–61.
- Caragiannis, I., and Procaccia, A. D. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175(9–10):1655–1671.
- Caragiannis, I.; Lai, J. K.; and Procaccia, A. D. 2011. Towards more expressive cake cutting. In *Proc. of 22nd IJCAI*, 127–132.
- Chechlárová, K., and Pillárová, E. 2012. On the computability of equitable divisions. *Discrete Optimization* 9(4):249–257.
- Chen, Y.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2013. Truth, justice, and cake cutting. *Games and Economic Behavior* 77:284–297. Preliminary version in AAI 10.
- Chevalyre, Y.; Dunne, P. E.; Endriss, U.; Lang, J.; Lemaître, M.; Maudet, N.; Padget, J.; Phelps, S.; Rodríguez-Aguilar, J. A.; and Sousa, P. 2006. Issues in multiagent resource allocation. *Informatica* 30:3–31.
- Cohler, Y. J.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2011. Optimal envy-free cake cutting. In *Proc. of 25th AAI*, 626–631.
- Deng, X.; Qi, Q.; and Saberi, A. 2012. Algorithmic solutions for envy-free cake cutting. *Operations Research* 60(6):1461–1476.
- Edmonds, J., and Pruhs, K. 2006a. Balanced allocations of cake. In *Proc. of 47th FOCS*, 623–634.
- Edmonds, J., and Pruhs, K. 2006b. Cake cutting really is not a piece of cake. In *Proc. of 17th SODA*, 271–278.
- Even, S., and Paz, A. 1984. A note on cake-cutting. *Discrete Applied Mathematics* 7:285–296.
- Gutman, A., and Nisan, N. 2012. Fair allocation without trade. In *Proc. of 11th AAMAS*, 719–728.
- Kash, I.; Procaccia, A. D.; and Shah, N. 2013. No agent left behind: Dynamic fair division of multiple resources. In *Proc. of 12th AAMAS*, 351–358.
- Kremer, I.; Nisan, N.; and Ron, D. 1999. On randomized one-round communication complexity. *Computational Complexity* 8(1):21–49.
- Kurokawa, D.; Lai, J. K.; and Procaccia, A. D. 2013. How to cut a cake before the party ends. In *Proc. of 27th AAI*, 555–561.
- Kushilevitz, E., and Nisan, N. 1996. *Communication Complexity*. Cambridge University Press.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proc. of 6th EC*, 125–131.
- Manabe, Y., and Okamoto, T. 2012. A cryptographic moving-knife cake-cutting protocol. In *Proc. of 2nd IWIGP*, 15–23.
- Muthukrishnan, S. 2005. *Data Streams: Algorithms and Applications*. Now Publishers.
- Procaccia, A. D. 2009. Thou shalt covet thy neighbor’s cake. In *Proc. of 21st IJCAI*, 239–244.
- Procaccia, A. D. 2013. Cake cutting: Not just child’s play. *Communications of the ACM* 56(7):78–87.
- Robertson, J. M., and Webb, W. A. 1998. *Cake Cutting Algorithms: Be Fair If You Can*. A. K. Peters.
- Woeginger, G. J., and Sgall, J. 2007. On the complexity of cake cutting. *Discrete Optimization* 4:213–220.