

How Good are Optimal Cake Divisions?

Steven J. Brams
New York University
steven.brams@nyu.edu

Michal Feldman
Hebrew University and Harvard University
mfeldman@huji.ac.il

John K. Lai
Harvard University
jklai@post.harvard.edu

Jamie Morgenstern
Carnegie Mellon University
jamiemmt@cs.cmu.edu

Ariel D. Procaccia
Carnegie Mellon University
arielpro@cs.cmu.edu

Abstract

We consider the problem of selecting fair divisions of a heterogeneous divisible good among a set of agents. Recent work (Cohler et al., AAAI 2011) focused on designing algorithms for computing optimal—social welfare maximizing (maxsum)—allocations under the fairness notion of envy-freeness. Maxsum allocations can also be found under alternative notions such as equitability. In this paper, we ask: how good are these allocations? In particular, we provide conditions for when maxsum envy-free or equitable allocations are Pareto optimal and give examples where fairness with Pareto optimality is not possible. We also prove that maxsum envy-free allocations have weakly greater welfare than maxsum equitable allocations when agents have structured valuations, and we derive an approximate version of this inequality for general valuations.

1 Introduction

How does one fairly divide a cake? This question has long been studied by mathematicians, economists, and political scientists (Brams and Taylor 1996; Robertson and Webb 1998), who view it as both a mathematical challenge and a metaphor for prominent real-world problems that involve the division of a divisible good. Such problems arise in the context of, e.g., land disputes and divorce settlements. In recent years, the rigorous study of cake cutting has gained significant traction within the AI community (Procaccia 2009; Chen et al. 2010; Caragiannis, Lai, and Procaccia 2011; Cohler et al. 2011; Walsh 2011; Zivan 2011), in part because it is seen as an important ingredient in the design of superior multiagent resource allocation methods (Chevalyre et al. 2006).

Most of the cake cutting literature focuses on the design of algorithms that compute fair cake divisions, under different interpretations of fairness. The notion of *envy-freeness* (EF) is perhaps the most prominent interpretation of fairness; an allocation is EF if each agent weakly prefers its own piece to the piece of cake allocated to any other agent. Note that the cake is usually heterogeneous (players prefer some parts of the cake to others, and different players may assign different values to the same piece of cake), so simply allocating pieces of equal size is not sufficient to guarantee EF. The notion of

equitability offers an alternative, incomparable, interpretation of fairness; an allocation is *equitable* (EQ) if all agents assign the same value to their own pieces. The existence of allocations satisfying these two notions of fairness is guaranteed under very mild assumptions and, in fact, both can be satisfied simultaneously (Alon 1987).¹

Fixing one of the two fairness criteria, cake cutting algorithms identify fair allocations. However, in general, multiple fair allocations exist, and some may be “better” than others (in ways to be specified). Recent work by Cohler et al. (2011) addresses this issue by adding an optimization objective. Specifically, they wish to maximize the (utilitarian) *social welfare*, that is, the sum of values the agents assign to their allocated pieces. Cohler et al. design algorithms that compute a *maxsum* (i.e., social-welfare-maximizing) allocation among all EF allocations. Their techniques can also be leveraged to compute a maxsum EQ allocation.

Intuitively, an overall maxsum EF (resp., EQ) allocation is superior to an arbitrary EF (resp., EQ) allocation. Nevertheless, we do not know *how good* maxsum EF or EQ allocations are; can one argue that they are truly more desirable than other allocations? Moreover, there are two notions of fairness to choose from; under which notion should one optimize social welfare?

Our approach and results. In economics, the quality of an allocation is often determined (in a binary fashion) via the criterion of *Pareto optimality* (PO): an allocation is PO if there is no *Pareto-dominating* allocation that gives at least as much value to all agents, and strictly more value to at least one agent. Note that a maxsum allocation is always PO, because a Pareto-dominating allocation would have higher social welfare. However, it is *a priori* unclear whether a maxsum EF or EQ allocation is PO among all possible allocations. Indeed, the answer depends on the notion of fairness.

¹*Proportionality* is another notion of fairness; a proportional division of a cake is one where the value of each of the n participating agents for its piece of cake is at least $1/n$ of its value for the entire cake. This relatively weak property is implied by EF (if the entire cake is allocated) and is not studied in this paper. We note here that EQ and social welfare (as defined in the next paragraph) assume interpersonal comparisons of utility among players, whereas EF does not, because agents compare their utility with the utility they would obtain from the pieces of the other players (instead of their utility with other players’ attributions of utility).

We first observe that, if there are only two agents, PO is guaranteed for maxsum EF allocations, maxsum EQ allocations, and even maxsum EF and EQ allocations (i.e., allocations that are maxsum among allocations that are both EF and EQ).

Our other results are more subtle and hinge on the structure of agents' valuation functions. As in previous papers (Chen et al. 2010; Cohler et al. 2011), we consider the special classes of *piecewise uniform valuations*, under which agents are simply interested in receiving as large a fraction as possible of a desired piece of cake; and the more general class of *piecewise constant valuations*, under which agents uniformly value certain pieces of cake. We show that under piecewise uniform valuations, maxsum EF allocations are always PO whereas there are cases where all maxsum EQ and maxsum EF+EQ allocations are not PO. Under piecewise constant valuations, there are examples with three agents such that all maxsum EF allocations are also not PO.


A second challenge we address compares the social welfare under maxsum EF and maxsum EQ allocations. We show that under piecewise constant valuations the social welfare of a maxsum EF allocation is at least as great as the social welfare of a maxsum EQ allocation. We also extend this result to general valuation functions albeit only approximately, in that (i) we optimize among allocations that are EF up to ϵ , and (ii) the inequality holds up to ϵ .

2 Preliminaries


A cake is represented by the interval $[0, 1]$, and a piece of cake X is a finite union of disjoint subintervals. There is also a set of agents $N = \{1, \dots, n\}$. The preferences of the agents over the cake are represented via *valuation functions* V_i , that map a given piece of cake to its value for the agent. The value is calculated as the integral of a non-negative Riemann integrable *value density function*, denoted by v_i . Formally, an agent's value $V_i(X)$ for a piece of cake X is given by $\sum_{I \in X} \int_I v_i(x) dx$. This definition guarantees that the agent valuations 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., $V_i([x, x]) = 0$. Non-atomicity means that we do not have to pay special attention to the endpoints of intervals; we can therefore treat open and closed intervals as equivalent. We also assume that agents' valuation functions are normalized so that the entire cake gives each agent value 1, that is, for all $i \in N$, $V_i([0, 1]) = 1$, or equivalently, $\int_0^1 v_i(x) dx = 1$.

While some of the cake-cutting literature assumes that valuations are *absolutely continuous* (see e.g., Brams, Jones, and Klamler 2012), i.e., that if any agent attaches zero value to a portion of the cake, then all other players do, the current paper does not employ this assumption.

Most of our results assume that the valuation functions have a specific structure. We say that a valuation function is *piecewise constant* if its corresponding value density function is piecewise constant, i.e., if the cake can be partitioned into a finite number of subintervals such that the density function is a constant function on each subinterval. An additional restriction is imposed by *piecewise uniform* valuation functions: on each subinterval, the density function is either



(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.

zero or some constant c_i , where the constant c_i is the same across different intervals. See Figure 1 for an illustration.

Piecewise constant and piecewise uniform valuation functions were the focus of several recent papers on cake cutting (Chen et al. 2010; Cohler et al. 2011). Piecewise uniform valuations have a natural interpretation: players have a desired piece of cake, and they value this piece uniformly, in the sense that they wish to receive as large a portion as possible of their desired piece. This is realistic, for example, when the cake represents access time to a shared backup server, and agents require as much time as possible but only when their computers are idle. Piecewise constant valuations are misleadingly simple, but they are in fact quite powerful; as we see in Section 5, general valuation functions can be approximated to arbitrary precision by piecewise constant valuation functions.

An *allocation* $A = (A_1, \dots, A_n)$ is an assignment of a piece of cake A_i to each agent i such that the pieces A_1, \dots, A_n are disjoint.²

We wish to focus on allocations that are fair; we consider two well-known notions of fairness. Given V_1, \dots, V_n , an allocation is *envy-free (EF)* if $V_i(A_i) \geq V_i(A_j)$ for all $i, j \in N$, and *equitable (EQ)* if $V_i(A_i) = V_j(A_j)$ for all $i, j \in N$. Envy-freeness guarantees that no agent wants to swap the piece that it is given with any other agent. Equitability ensures that each agent obtains the same value for its piece as all other agents obtain for their pieces.

A third criterion for allocations will help us gauge their quality. We say that an allocation $A = (A_1, \dots, A_n)$ is *Pareto dominated* by another allocation $A' = (A'_1, \dots, A'_n)$ if $V_i(A'_i) \geq V_i(A_i)$ for all $i \in N$, and there exists $i \in N$ such that $V_i(A'_i) > V_i(A_i)$. An allocation is *Pareto optimal (PO)* if it is not Pareto dominated by any other allocation.

²Note that this definition allows for intervals to be discarded.

Under normalization (which implicitly implies an interpersonal comparison of utility) it is meaningful to consider the sum of the agent valuations in a given allocation. The (utilitarian) *social welfare* of an allocation A is given by $sw(A) = \sum_{i=1}^n V_i(A_i)$. An allocation A is *maxsum* among a set of possible allocations \mathcal{S} if $sw(A) = \max_{A' \in \mathcal{S}} sw(A')$. In particular, we shall be interested in the properties of the maxsum allocation when \mathcal{S} is the set of EF allocations, EQ allocations, and allocations that are both EF and EQ. These allocations will be referred to as maxsum EF, maxsum EQ, and maxsum EF+EQ allocations, respectively.

Throughout the paper, some proofs are only sketched or omitted due to space limitations.

3 Pareto Optimality of Maxsum Allocations

In this section, we study the Pareto optimality of maxsum allocations. In particular, we establish the Pareto optimality of maxsum EF, EQ, and EF+EQ allocations in the case of two agents and general valuations, and complement this result by showing that for three agents or more, these allocations are not necessarily Pareto optimal.

Theorem 1. *For general valuations and two agents, every maxsum EF, EQ, or EF+EQ allocation is PO.*

Note that the two-agent case has special significance (for example, in the context of divorce settlements), and indeed the main result of Caragiannis et al. (2011) captures only the two agent case.

Before proving Theorem 1, we introduce the notion of ratio-based allocations for the two-agent setting. These allocations have been used by Cohler et al. (2011) to find maxsum EF allocations.

For a given pair of valuation densities v_1, v_2 , let $Y_{i \text{ op } j} = \{x : v_i(x) \text{ op } v_j(x)\}$. For instance, $Y_{1 \geq 2}$ gives all intervals where agent 1's value density function is weakly greater than agent 2's, and $Y_{1 > 2}$ gives all intervals where agent 1's value density function is strictly greater than agent 2's. Additionally, let Y_1, Y_2 denote intervals that are only desired by agent 1 and only desired by agent 2, respectively. Denote the ratio of the value density functions by $R_1(x) = v_1(x)/v_2(x)$ and $R_2(x) = v_2(x)/v_1(x)$. Let

$$Y_{R_1 \text{ op } r} = \{x : v_1(x) \leq v_2(x), v_2(x) > 0, R_1(x) \text{ op } r\}$$

$$Y_{R_2 \text{ op } r} = \{x : v_2(x) \leq v_1(x), v_1(x) > 0, R_2(x) \text{ op } r\},$$

where $\text{op} \in \{>, =\}$.

Definition 2. An allocation $A = (A_1, A_2)$ is *ratio-based* if $Y_1 \subseteq A_1, Y_2 \subseteq A_2$ and either one of the following holds:

- There exists an $r^* \in [0, 1]$ such that

$$A_1 = Y_{1 > 2} \cup Y_{R_1 > r^*} \cup C,$$

where $C \subseteq Y_{R_1 = r^*}$.

- There exists an $r^* \in [0, 1]$ such that

$$A_2 = Y_{2 > 1} \cup Y_{R_2 > r^*} \cup C,$$

where $C \subseteq Y_{R_2 = r^*}$.

We refer to agent 1 as the *receiving agent* in the first case and agent 2 as the receiving agent in the second case. We refer to r^* as the *critical ratio*.

In a ratio-based allocation, the receiving agent is always allocated intervals that it strictly desires, as well as some intervals weakly desired by the other agent. For the special case where the critical ratio is 1, both agents can be seen as receiving agents. In this case, the allocation is maxsum since all intervals are allocated to agents who weakly desire the interval. When the critical ratio is less than 1, there is a unique receiving agent i that receives all intervals it weakly desires ($Y_{i > 3-i}$) along with some intervals strictly desired by the other agent. This necessarily results in a loss of welfare relative to the maxsum allocation. However, ratio-based allocations minimize the obtained loss. This is formalized in the following lemma.

Lemma 3. *Let $A = (A_1, A_2)$ be a ratio-based allocation with agent 1 as the receiving agent such that $v = V_1(A_1) \geq V_1(Y_{1 \geq 2})$. It holds that:*

1. *For every allocation $A' = (A'_1, A'_2)$ such that $V_1(A'_1) = v$, $sw(A) \geq sw(A')$.*
2. *For every allocation $A' = (A'_1, A'_2)$ such that $V_1(A'_1) > v$, $sw(A) > sw(A')$.*

The analogous assertion holds for agent 2.

Proof sketch. The proof of the lemma closely resembles the proof of Theorem 3 in Cohler et al. (2011). Among all allocations that grant agent 1 value v , the allocation that maximizes welfare is one in which agent 1 is first allocated all the intervals it strictly desires, and then, possibly, intervals that are strictly desired by agent 2, in a decreasing order of $R_i(x)$. The first part entails no loss in welfare. The second part may entail some loss, but allocating these intervals in a decreasing order of $R_i(x)$ ensures that this is the lowest possible loss. In addition, if agent 1 receives value greater than v , it must come from additional intervals that are strictly desired by agent 2. This entails a greater loss in welfare. \square

The following is an immediate corollary of Lemma 3.

Lemma 4. *Every ratio-based allocation is PO.*

We require one additional lemma for the proof of Theorem 1.

Lemma 5. *Every maxsum EF allocation allocates all intervals that are desired by some agent.*

Proof sketch. Any allocation that discards intervals that are desired by some agent can be concatenated by an EF division of the discarded interval, maintaining EF while increasing social welfare. The discarded interval can be allocated, e.g., through the cut-and-choose method, where one agent splits the interval into equi-value pieces (according to its valuation) and the other agent chooses the preferred piece. \square

Proof of Theorem 1. We address the three different allocation types.

Maxsum EF. Cohler et al. (2011) establish the existence of a ratio-based maxsum EF allocation.³ We distinguish between two cases, as follows. If there exists an EF alloca-

³While they focus on piecewise linear valuation functions, their proof holds for general valuation functions.

tion that is maxsum among all allocations, then every maxsum EF allocation is trivially PO. Otherwise, it is shown by Cohler et al. (2011) that there must exist an agent i such that $V_i(Y_{i \geq 3-i}) < 1/2$. Wlog, suppose $V_1(Y_{1 \geq 2}) < 1/2$. Cohler et al. (2011) establish the existence of a ratio-based allocation that gives agent 1 value of exactly $1/2$ and is maxsum EF. Let $A = (A_1, A_2)$ be such an allocation. By Lemma 4, allocation A is PO. Let $A' = (A'_1, A'_2)$ be another maxsum EF allocation. In what follows we show that A' is PO. We distinguish between three cases.

1. If $V_1(A'_1) = 1/2$, then, since A' is maxsum EF, it follows that $V_2(A'_2) = V_2(A_2)$. In this case, the fact that A is PO implies that A' is PO as well.
2. If $V_1(A'_1) < 1/2$, then $V_1(A'_2) < 1/2$ (otherwise, contradicting EF). It follows by Lemma 5 that A' is not maxsum EF, a contradiction.
3. If $V_1(A'_1) > 1/2$, then we get $V_1(A'_1) > 1/2 > V_1(Y_{1 \geq 2})$. It follows by Lemma 3 that $sw(A') < sw(A)$, in contradiction to A' being a maxsum EF allocation.

Maxsum EQ. We distinguish between two cases.

- $V_1(Y_{1 \geq 2}) \geq V_2(Y_{2 > 1})$ and $V_2(Y_{2 \geq 1}) \geq V_1(Y_{1 > 2})$. In this case we show that there exists a maxsum EQ allocation that is maxsum among all allocations. This, in turn, implies that every maxsum EQ allocation is PO. In particular, allocate $Y_{1 > 2}$ to agent 1, $Y_{2 > 1}$ to agent 2, and split $Y_{1=2}$ such that the agents' values for their pieces are equal. To see why this is feasible, note that if we give $Y_{1=2}$ to agent 1 in its entirety, then agent 1 has a greater value. On the other hand, if we give all of $Y_{1=2}$ to agent 2, then agent 2 has a greater value. Therefore, there must exist some allocation of $Y_{1=2}$ that equalizes their values. This allocation is maxsum among all allocations.
- Wlog, suppose $V_1(Y_{1 \geq 2}) < V_2(Y_{2 > 1})$. We claim that in this case there exists a ratio-based allocation with agent 1 as the receiving agent that is EQ. To see this, note that as the critical ratio decreases from 1 to 0, agent 1 goes from receiving all of $Y_{1 \geq 2}$ to receiving the entire cake, i.e., from a value of $V_1(Y_{1 \geq 2})$ to a value of 1. On the other hand, agent 2 goes from receiving all of $Y_{2 > 1}$ to receiving none of the cake, i.e., from value $V_2(Y_{2 > 1}) > V_1(Y_{1 \geq 2})$ to 0. Therefore, the agents' values must cross at some point, and the assertion follows. By Lemma 4 this allocation is PO, and hence maxsum EQ. Clearly, any maxsum EQ allocation must grant each agent the same value as in the ratio-based maxsum EQ allocation. It follows that every maxsum EQ allocation is PO.

Maxsum EF+EQ. In every maxsum EQ allocation, both agents receive value at least $1/2$. If this were not true, then the agents could swap allocations and obtain a maxsum EQ allocation with greater social welfare. Since both agents receive value at least $1/2$, the maxsum EQ allocation is also EF. It follows that for two agents, the set of maxsum EF+EQ allocations coincides with the set of maxsum EQ allocations, for which the assertion of the theorem is proved above. \square

We next turn to investigate maxsum EF, maxsum EQ, and maxsum EF+EQ allocations under restricted valuations. As it turns out, at least under piecewise uniform valuation functions, maxsum EF allocations are always PO whereas maxsum EQ and maxsum EF+EQ allocations may not be.

Theorem 6. *For piecewise uniform valuations, every maxsum EF allocation is PO.*

Proof sketch. When agent valuations are piecewise uniform, a sufficient condition for PO is that all intervals desired by at least one agent are allocated to an agent that has positive density on the entire interval. To see why this is true, recall that when agents have piecewise uniform valuations, their total value is exactly determined by the total length of desired intervals they receive. If all desired intervals are allocated to agents with positive density, then a Pareto-dominant allocation cannot exist because this would require additional desired lengths to be created. It remains to show that a maxsum EF allocation must have this property.

Suppose that a maxsum EF allocation $A = (A_1, \dots, A_n)$ allocates some intervals to agents that do not desire them or discards intervals altogether. Let X' denote these intervals. Under piecewise uniform valuations, we can split X' into subintervals on which agent densities are constant, and then give each agent a $1/n$ share of each of these subintervals. We can append this allocation of X' to A . Envy is not created, because each agent i has value exactly $(1/n)V_i(X')$ for every piece in this allocation, but social welfare increases, contradicting the assumption that A is maxsum. \square

Theorem 7. *For piecewise uniform valuations and three agents, there are valuation functions where all maxsum EQ and EF+EQ allocations are not PO.*

Proof. Consider the following valuations. Agents 1 and 2 desire $[0, 0.1]$ and agent 3 desires all of $[0, 1]$. A maxsum EQ or maxsum EQ+EF allocation must split $[0, 0.1]$ between agents 1 and 2 and allocate $[0, 1]$ to agent 3 so that agent 3 receives value exactly 0.5. This is not PO because we can split $[0, 0.1]$ between agents 1 and 2 and give agent 3 all of $[0.1, 1]$. \square

While there are cases where no maxsum EQ or EF+EQ allocation is PO under piecewise uniform valuations, we need to move to piecewise constant valuations in order to find cases where no maxsum EF allocation is PO.

Theorem 8. *For piecewise constant valuations and three agents, there are cases where no maxsum EF allocation is PO.*

We view Theorem 8, whose proof is omitted for space, as one of our main results because of the significance of maxsum EF allocations (Cohler et al. 2011). Finding an initial example required automated search, and proving that (a modified version of) this example admits a unique maxsum EF allocation relies on reasoning about the linear programming formulation of the problem.

4 Maxsum EQ vs. Maxsum EF Allocations

In this section, we show that for piecewise linear valuations, a maxsum EF allocation has social welfare at least as large as any maxsum EQ allocation. First, we show the theorem for piecewise constant valuations. We obtain an approximate version of this result for general valuation functions.

Denote the social welfare of a maxsum EF (resp., EQ) allocation by OPT_{EF} (resp., OPT_{EQ}). Note that the two-agent version of the inequality $\text{OPT}_{\text{EQ}} \leq \text{OPT}_{\text{EF}}$, for any valuation functions, follows from the fact that a maxsum EQ allocation is also EF, which was established in passing in the proof of Theorem 1. As a recap, both agents receive value at least $1/2$ in a maxsum EQ allocation, and for two agents, this is a sufficient condition for envy-freeness.

For three agents, this argument no longer holds, even in the case of piecewise constant valuations: a maxsum EQ allocation must give utility at least $1/3$ to each agent, but this does not imply EF.

For example, consider the piecewise uniform valuations where agents 1 and 2 value the whole cake (with density 1) and agent 3 only values $[0.8, 1]$ (with density 5). A maxsum EQ allocation would be to give agent 1 $[0, 5/11]$, agent 2 $[5/11, 10/11]$, and agent 3 $[10/11, 1]$. Each agent receives value $5/11$, yet agent 3 envies agent 2.

Another interesting (but common) feature of this example is that $\text{OPT}_{\text{EQ}} < \text{OPT}_{\text{EF}}$, with a strict inequality. One EF allocation is to give $[0.8, 1]$ to agent 3 and split $[0, 0.8]$ between agents 1 and 2. This has social welfare of 1.8 compared to the maxsum EQ welfare of $15/11 \approx 1.364$.

Having built some intuition, we next present the main result of this section. An ϵ -EF allocation is one where $V_i(A_i) \geq V_i(A_j) - \epsilon$ for all $i, j \in N$. Let $\text{OPT}_{\epsilon\text{-EF}}$ denote the social welfare under a maxsum ϵ -EF allocation.

Theorem 9. *For piecewise linear valuations,*

$$\text{OPT}_{\text{EQ}} \leq \text{OPT}_{\text{EF}}.$$

Moreover, for general valuation functions and any $\epsilon > 0$,

$$\text{OPT}_{\text{EQ}} \leq \text{OPT}_{\epsilon\text{-EF}} + \epsilon.$$

The proof of Theorem 9 relies on a connection between piecewise constant valuation functions and market equilibria for a collection of divisible goods inspired by the work of Reijnerse and Potters (1998). Before we begin the proof, we draw this connection and cite the relevant results from the market equilibria literature required in the proof.

A *linear Fisher market* is a market where agents $N = \{1, \dots, n\}$ have additive, linear utility functions for a set $G = \{1, \dots, m\}$ of divisible goods. Each agent $i \in N$ is given a budget e_i and has a utility u_{ij} for each good $j \in G$. A *feasible allocation* gives a fraction x_{ij} of good j to agent i such that no good is over-allocated. The agent's total utility from an allocation x_{ij} is $\sum_j u_{ij}x_{ij}$. When agent valuations are piecewise constant, utilities in a feasible Fisher market allocation can be replicated in the cake cutting setting.

Lemma 10. *Let A_1, \dots, A_n be an allocation in the cake cutting setting. Define a Fisher market with the same agents, and a good j corresponding to each A_j and $u_{ij} = V_i(A_j)$.*

Let x_{ij} be a feasible allocation of goods in the Fisher market. Suppose the agents' valuations are piecewise constant. Then, there exists an allocation A'_1, \dots, A'_n such that $V_i(A'_i) = \sum_j u_{ij}x_{ij}$. In other words, we can replicate agent utilities in the Fisher market with an allocation of the cake.

Proof. Given a feasible allocation x_{ij} in the Fisher market, create an allocation A'_1, \dots, A'_n as follows. For each original piece A_j , split A_j into subintervals on which every agent's value density function is constant (this is possible since agent value densities are piecewise constant). Give agent i a fraction x_{ij} of each subinterval. Agent i 's value from its share of A_j is $x_{ij}V_i(A_j) = x_{ij}u_{ij}$, and summing over all intervals proves the lemma. \square

Corollary 11. *Lemma 10 also holds for piecewise linear valuations.*

Proof. Given a feasible allocation x_{ij} in the Fisher market, create an allocation A'_1, \dots, A'_n as follows. For each original piece A_j , split A_j into subintervals on which every agent's value density function is linear. Allocate each of these subintervals, in the following way. Give agent i two pieces of A_j , each of which is a $\frac{x_{ij}}{2}$ fraction of A_j , one starting at the leftmost unallocated part of A_j , one starting at the rightmost unallocated part of A_j . Agent i 's value from its share of A_j is still $x_{ij}V_i(A_j) = x_{ij}u_{ij}$, and agent i 's value from agent k 's share of A_j is $x_{kj}V_i(A_j) = x_{kj}u_{ij}$, as required. \square

Linear Fisher markets have the following very special properties (see e.g., Vazirani 2007).

Theorem 12. *Consider a linear Fisher market where agent i has budget e_i , $\sum_{i \in N} e_i = 1$, and each good gives at least one agent positive utility. There exists a price vector $p = (p_1, \dots, p_{|G|})$, $p_j > 0$, $\sum_{j \in G} p_j = 1$, and a feasible allocation x_{ij} such that:*

1. $\forall j \in G, \sum_{i \in N} x_{ij} = 1$,
2. $\forall i \in N, j \in G$, If $x_{ij} > 0$, then $j \in \text{argmax}_{j'} (u_{ij'}/p_{j'})$,
3. $\forall i \in N, \sum_{j \in G} p_j x_{ij} = e_i$.

Leveraging this result, we prove Theorem 9. We first show the result for piecewise constant valuations. Next, we approximate the piecewise linear valuations of the players with piecewise constant valuations to prove the theorem.

Proof of Theorem 9. We first show the result for piecewise constant valuation functions. Begin with the maxsum EQ allocation $A^* = (A_1^*, \dots, A_n^*)$. Construct a Fisher market where good j corresponds to A_j^* , $u_{ij} = V_i(A_j^*)$ and each agent has budget $e_i = 1/n$. Let p, x_{ij} be the price vector and feasible allocation guaranteed by Theorem 12. Consider the allocation A'_1, \dots, A'_n described in Lemma 10. We need to show that this allocation is EF and yields total welfare weakly greater than that of the original maxsum EQ allocation. Due to Lemma 10, we can relate the values for A'_1, \dots, A'_n to the utilities in the Fisher market.

The proof that A'_1, \dots, A'_n is EF appears in Reijnerse and Potters (1998); the next equation replicates it for completeness. Let $u_i^* = \max_k (u_{ik}/p_k)$.

$$\begin{aligned} V_i(A'_i) &= \sum_k u_{ik} x_{ik} = \sum_k \frac{u_{ik}}{p_k} p_k x_{ik} \\ &= \sum_k u_i^* p_k x_{ik} = u_i^*/n \\ V_i(A'_j) &= \sum_k u_{ik} x_{jk} = \sum_k \frac{u_{ik}}{p_k} p_k x_{jk} \\ &\leq \sum_k u_i^* p_k x_{jk} = u_i^*/n \end{aligned}$$

It remains to show that $\sum_i V_i(A'_i) \geq \sum_i V_i(A_i^*)$. Suppose $V_i(A_i^*) = C$ for all $i \in N$; then $\text{OPT}_{\text{EQ}} = \sum_i V_i(A_i^*) = nC$. u_i^* maximizes u_{ik}/p_k , so u_i^* is at least u_{ii}/p_i , the utility to price ratio for the good in the Fisher market corresponding to A_i^* . Therefore, $V_i(A'_i) = u_i^*/n \geq u_{ii}/(np_i) = C/(np_i)$.

Then,

$$\text{OPT}_{\text{EF}} \geq \sum_i V_i(A'_i) \geq \sum_i \frac{C}{np_i} = \frac{C}{n} \sum_i \frac{1}{p_i}.$$

Since $\sum_i p_i = 1$, $\sum_i (1/p_i)$ is minimized by $p_i = 1/n$ for each i and is at least n^2 . Therefore,

$$\text{OPT}_{\text{EF}} \geq \frac{C}{n} \sum_i \frac{1}{p_i} \geq \frac{C}{n} n^2 = nC = \text{OPT}_{\text{EQ}}.$$

Now, leveraging this result, we prove the theorem for piecewise linear valuations. Partition $[0, 1]$ such that on every subinterval, each player's valuation is linear. Call this partition $P = (p_1, \dots, p_k)$.

Approximate the valuations of each player on each $p_j = [a, b]$ with the constant valuation for all $x \in [a, b]$

$$v'_i(x) = v_i \left(\frac{b-a}{2} \right)$$

(e.g. approximate the valuations with their midpoint values on each of the intervals p_i). Notice that

$$V'_i([a, b]) = V_i([a, b])$$

Construct a Fisher market as in the linear case, with $u_{ij} = V'_i(A_j^*) = V_i(A_j^*)$. Then, construct an envy-free allocation X' given by the proof of Corollary 11. Thus, the valuations will be exact on this allocation, and from the previous proof about piecewise constant valuations, the allocation is envy-free and has value at least that of X .

Next we establish our result for general valuation functions V_1, \dots, V_n (with Riemann integrable value density functions). For $\epsilon > 0$, Riemann integrability of v_1, \dots, v_n implies that for all $i \in N$ there are $0 = x_1 < \dots < x_m = 1$ such that the upper Darboux sum of v_i satisfies

$$\begin{aligned} 1 &= \int_{x=0}^1 v_i(x) dx \\ &\leq \sum_{k=1}^m \left[(x_k - x_{k-1}) \cdot \left(\sup_{x \in [x_{k-1}, x_k]} v_i(x) \right) \right] \leq 1 + \frac{\epsilon}{n}. \end{aligned} \tag{1}$$

For every $k = 1, \dots, m$ and every $y \in [x_{k-1}, x_k]$, let $v'_i(y) = \sup_{x \in [x_{k-1}, x_k]} v_i(x)$. We claim that the corresponding piecewise constant valuation functions V'_1, \dots, V'_n approximate the original valuation functions in the sense that for every piece of cake X ,⁴

$$V_i(X) \leq V'_i(X) \leq V_i(X) + \frac{\epsilon}{n}. \tag{2}$$

Indeed, the left hand side of the inequality is trivial, and the right hand side follows from Equation (1) and the fact that $v'_i(x) \geq v_i(x)$ for all $x \in [0, 1]$:

$$\begin{aligned} V'_i(X) - V_i(X) &= \int_X (v'_i(x) - v_i(x)) dx \\ &\leq \int_{x=0}^1 (v'_i(x) - v_i(x)) dx \leq \frac{\epsilon}{n}. \end{aligned}$$

Assume as before that the maxsum EQ allocation A^* satisfies $V_i(A_i^*) = C$ for all $i \in N$. It therefore holds that $V'_i(A_i^*) \geq C$ for all $i \in N$. Using the same arguments as before, there exists an allocation A' that is EF with respect to V'_1, \dots, V'_n and satisfies

$$\sum_{i \in N} V'_i(A'_i) \geq nC = \sum_{i \in N} V_i(A_i^*) = \text{OPT}_{\text{EQ}}.$$

Equation (2) directly implies that the allocation A' is ϵ -EF (in fact, (ϵ/n) -EF) with respect to the valuations V_1, \dots, V_n . Therefore, it holds that

$$\begin{aligned} \text{OPT}_{\epsilon\text{-EF}} &\geq \sum_{i \in N} V_i(A'_i) \geq \sum_{i \in N} \left(V'_i(A'_i) - \frac{\epsilon}{n} \right) \\ &= \sum_{i \in N} V'_i(A'_i) - \sum_{i \in N} \frac{\epsilon}{n} \geq \text{OPT}_{\text{EQ}} - \epsilon. \end{aligned}$$

□

5 Discussion

Our work can be seen as another step on the path to identifying the most desirable allocations of divisible goods. In recent work, Brams et al. (2012) coined the term *perfect allocations* to describe allocations that are PO, EF, and EQ. Unfortunately, they show that such allocations may not exist when there are three or more agents, however many cuts are allowed. We therefore argue that maximizing social welfare under a subset of these three properties provides an especially appealing solution, but as we discuss below, there are trade-offs between the different properties.

One may wonder, in light of Theorem 9, whether a maxsum EF allocation is superior to a maxsum EQ allocation. While we believe that this is often true, we wish to add a caveat. Consider an example where there are three agents with value density functions $v_1(x) = v_2(x) = 1$, $v_3(x) = 2x$. A maxsum EF allocation gives $[0, 1/3]$ to agent 1, $[1/3, 2/3]$ to agent 2, and $[2/3, 1]$ to agent 3, for a sum of $1/3 + 1/3 + 5/9 \approx 1.22$. This allocation also happens to be PO. But there is a maxsum EQ allocation that is also EF (by dividing the left portion of the cake between agents 1 and 2

⁴It may be the case that $V'_i([0, 1]) > 1$.

in a way that 3 does not envy either) and gives each agent a value of roughly 0.39, for a slightly lower sum of 1.17. The latter allocation seems more desirable, because it maximizes the minimum value to the agents. Indeed, the EF allocation creates significant inequity between agents 1 and 2, on the one hand, and agent 3 on the other ($1/3$ vs. $5/9$); this 67% difference in values in exchange for only a 4% higher social welfare, compared with EQ (1.22 vs. 1.17), arguably tips the balance in favor of the maxsum EQ allocation: it not only gives all agents the same “fair share,” unlike the maxsum EF allocation, but it is also EF.

We have shown that maxsum EF allocations may not be PO, and hence one may consider choosing an EF allocation that Pareto-dominates the maxsum EF allocation. However, in the examples that we have been able to construct where the maxsum EF allocation is indeed not PO, the difference in social welfare between the maxsum EF allocation and its Pareto-dominating allocation is very small. Bounding this difference (or ratio) remains an open question (which is somewhat related to work on the so-called *price of fairness* (Caragiannis et al. 2009)), but if it is indeed always small, we would argue that preserving EF is more important than a small gain in social welfare.

Another alternative is forcing the allocation to satisfy PO by taking the maxsum over both EF and PO. Reijniere and Potters (1998) designed an elaborate algorithm that computes EF and PO allocations. However, these allocations are not maxsum necessarily. The techniques of Cohler et al. (2011) enable the computation of maxsum EF allocations, which are not necessarily PO. Our most important, and presumably quite challenging, open problem is therefore the construction of a (tractable) algorithm that computes maxsum EF and PO allocations.

Returning to our eponymous question, “how good are optimal cake divisions?”, we have shown that maxsum cake divisions are imperfect, and we have crystallized some of the trade-offs between them. Our contributions inform the discussion of good methods for resource allocation by (i) ruling out the possibility that maxsum EF allocations are always superior to other allocations (by showing that they may not be PO), and (ii) demonstrating that moving from EF to the egalitarian notion of EQ can only decrease the utilitarian social welfare.

References

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

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

Brams, S. J.; Jones, M. A.; and Klamler, C. 2012. *N*-person cake-cutting: There may be no perfect division. *American Mathematical Monthly*, forthcoming.

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2009. The efficiency of fair division. In *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE)*, 475–482.

Caragiannis, I.; Lai, J. K.; and Procaccia, A. D. 2011. Towards more expressive cake cutting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 127–132.

Chen, Y.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2010. Truth, justice, and cake cutting. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 756–761.

Chevalere, 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 *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)*, 626–631.

Procaccia, A. D. 2009. Thou shalt covet thy neighbor’s cake. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 239–244.

Reijniere, J., and Potters, J. A. M. 1998. On finding an envy-free Pareto-optimal division. *Mathematical Programming* 83:291–311.

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

Vazirani, V. V. 2007. Combinatorial algorithms for market equilibria. In Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V., eds., *Algorithmic Game Theory*. Cambridge University Press. chapter 5.

Walsh, T. 2011. Online cake cutting. In *Proceedings of the 2nd International Conference on Algorithmic Decision Theory (ADT)*, 292–305.

Zivan, R. 2011. Can trust increase the efficiency of cake cutting algorithms? In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 1145–1146.