A Tradeoff between Fairness and Efficiency in Cake-cutting
– DRAFT –

Erel Segal-Halevi

Bar-Ilan University, Ramat-Gan 5290002, Israel,
erelsgl@gmail.com

Abstract. Global efficiency and individual fairness are two important considerations in resource allocation. Efficiency aims to improve the sum of values of all agents (also called the "utilitarian social welfare"). Fairness aims to ensure that every agent has a value above a certain threshold. These considerations are often conflicting and cannot be satisfied simultaneously. This raises the question of what compromise can be attained between them. In this paper we study this question in the context of the classic problem of dividing a heterogeneous resource ("cake") among several agents with different valuations. We show that, indeed, a viable compromise exists between the two criteria. It is possible to attain a constant-factor approximation to the maximum utilitarian welfare, while still guaranteeing a value of 1/O(n) to each agent. Our result is constructive: it is based on an algorithm taking an existing utilitarian-optimal cake allocation and modifying it in order to guarantee a sufficient value to each agent, without much reduction in the utilitarian welfare. Our algorithm can be interpreted as a procedure for *land reform* - redividing an already-divided land-estate in order to increase fairness.

1 Introduction

Fair division of land and other resources among agents with different preferences has been an important issue since Biblical times. Today it is an active area of research in the interface of computer science [16, 14] and economics [13]. Its applications range from politics [5, 4] to multi-agent systems [8].

A key challenge in fair division is determining the criteria by which a division should be selected. [19] introduced the elementary and most basic fairness criterion, now termed *proportionality*: each of the n agents should get a piece which he values as worth at least $1/n$ of the value of the entire cake. This criterion corresponds to the Rawlsian principle of justice [15], which says that social welfare is measured by the welfare of the poorest individual. The proportionality criterion indeed guarantees that even the poorest individual receives at least his fair share of $1/n$ the total value.

In addition to fairness, it is desired that an allocation be economically efficient. A common measure for economic efficiency is the sum of utilities of all agents. This sum is usually termed the *utilitarian social welfare* since it corresponds to the utilitarian principle [12], which says that social welfare is measured by the sum of welfare measures of all individuals.

Proportionality and utilitarianism are often conflicting. This is shown in the example below, a variant of which appears in several papers [3, 2, 6, 1]:

Example 1. Let k be an integer and $n = k^2$. A land-estate of area n has to be divided among n agents. The land-estate is divided to k regions, $R_1, \ldots, R_k$, with an area of k per region. There are two groups of agents:

- k *regional agents*: each agent i in this group wants a specific region R_i. The utility of agent i is k times the area it receives in region R_i (so the maximum possible utility per agent is $k^2 = n$).
- $n - k$ *dispersed agents*: For each agent in this group, the utility equals the area it receives (so the maximum possible utility is again n).

Note that the utility functions of the agents are normalized such that the maximum possible utility is the same for all agents.

The utilitarian-optimal allocation gives the entire region R_i to the regional agent i; the utilitarian welfare in this case is k^3. However, the dispersed agents receive nothing so the division is not proportional.

Any proportional allocation must give to each dispersed agent an area of at least 1, so the total area given to the dispersed agents is at least $n - k$. This leaves an area of at most k to give to the regional agents. Hence, the sum of utilities of the rich agents is at most k^2 and the total sum is at most $k^2 + (n - k) = 2k^2 - k$. The utilitarian welfare drops, relative to its maximum, by a factor of $k^2/(2k - 1) = n/(2\sqrt{n} - 1) = \Theta(\sqrt{n})$. □

In the above example, a utilitarian-optimal allocation gives a value of 0 to the less fortunate agents, while a proportional allocation causes a loss in utilitarian-welfare that goes to infinity as the number of agents grows. These pessimistic results raise the question of whether there is some compromise that attains reasonable levels of utilitarian welfare and proportionality even when the number of agents is large. Our first contribution is an affirmative answer to this question.

Theorem 1. *[Section 3] For every $r \in [0, 1]$ and every number n of agents, there exists an allocation which simultaneously attains a utilitarian welfare of $1/r$ the optimum, and guarantees each agent at least $(1 - r)/n$ of its total value.*

As an example, there exists a half-proportional allocation (giving each agent at least $1/(2n)$ its total cake value) that attains a utilitarian welfare of at least $1/2$ the optimum.

This result is encouraging. However, its disadvantage is that the protocol applied to attain this division gives each agent a large number of disconnected pieces (the number of pieces is n times the denominator of r, so the half-proportional allocation gives each agent $2n$ disconnected pieces). Is it possible to attain a similar result when the pieces must be connected?

Here our answer is negative - a constant-factor approximation to utilitarian welfare is impossible if we want to guarantee positive value to all agents:

Theorem 2. *[Section 4] For every number n of agents, there exist agent valuations such that, in every connected allocation that gives a positive value to each agent, the utilitarian welfare is at most $\Theta(\sqrt{\frac{\lg \lg n}{\lg n}})$ of the optimum.*

In light of this impossibility, the next question is: what reduction in the utilitarian welfare is sufficient for guaranteeing a positive value to all agents? Our next theorem provides an upper bound:

Theorem 3. *[Section 5] For every number n of agents, there exists an allocation which simultaneously approximates the optimal utilitarian welfare to a factor $1/O(\log n)$ and guarantees each agent a positive value.*

The positive-value guarantee is arguably quite weak (but in light of Theorem 2, it is far from trivial). We view it as a first step towards future results, which (hopefully) will guarantee each agent a constant-factor approximation to $1/n$ with logarithmic loss in utilitarian welfare.

1.1 Model

A *cake* C is a Borel subset of some Euclidean plane $\mathbb{R}^d$. In this paper we focus on cakes that are intervals ($d = 1$).

C has to be divided among $n \geq 1$ *agents*. Each agent $i \in \{1, \ldots, n\}$ has a *value-density* function v_i, which is an integrable, non-negative and bounded function on C. The *value* of a piece X to agent i is marked by $V_i(X)$ and it is the integral of its value-density:

$$V_i(X) = \int_{x \in X} v_i(x) dx$$

The V_i are measures and are *absolutely continuous with respect to the Lebesgue measure*, i.e., any piece with zero length has zero value to all agents. Therefore, we do not need to worry about who gets the end-points of a piece.

Unless otherwise stated, the agents' valuations are normalized such that for each agent, the value of the entire cake is 1.

The agents' valuations are their private information. A division protocol accesses the agents' valuations via *queries*. Standard cake-cutting protocols use two types of queries [16]: an *eval query* asks an agent to reveal its value for a specified piece of cake; a *mark query* asks an agent to mark a piece of cake with a specified value. As usual in the cake-cutting literature since [19], the fairness guarantees of our division protocols are valid for every agent answering the queries truthfully, regardless of the behavior of the other agents.

An *allocation* is a vector of n pieces, $X = (X_1, \ldots, X_n)$, one piece per agent, such that the X_i are pairwise-disjoint and $\cup_{i=1}^{n} X_i \subseteq C$. Note that some cake may remain unallocated, i.e, *free disposal* is assumed.

An allocation is called *connected* if for all i, each X_i is connected (in the one dimensional case, this means that all pieces are intervals).

An allocation is assessed by its *social welfare*. The social welfare of an allocation is a certain aggregate function of the normalized values of the agents (the normalized value is the piece value divided by the total cake value). In this paper we focus on two common social welfare functions: egalitarian and utilitarian [13]. We normalize them such that the maximum welfare is 1:

– *Egalitarian welfare* - the minimum of the agents' normalized values:

$$W_{egal}(X) = \min_{i\in\{1,\ldots,n\}} \frac{V_i(X_i)}{V_i(C)}$$

– *Utilitarian welfare* - the arithmetic mean of the agents' normalized values:

$$W_{util}(X) = \frac{1}{n} \sum_{i\in\{1,\ldots,n\}} \frac{V_i(X_i)}{V_i(C)}$$

A cake-allocation is called *utilitarian* if it maximizes the utilitarian welfare. Given a fraction $r \in [0, 1]$, a cake-allocation is called r*-utilitarian* if its utilitarian welfare is at least r times the maximum.

Egalitarian welfare is related to the proportionality criterion: a *proportional allocation* is an allocation with an egalitarian welfare of at least $1/n$ - every agent receives at least $1/n$ of the total cake value:

$$\forall i \in \{1, \ldots, n\} : V_i(X_i) \geq \frac{V_i(C)}{n}$$

Given a fraction $r \in [0, 1]$, we call an allocation r*-proportional* if its egalitarian welfare is at least r/n. We call an allocation *partially-proportional* if it is r-proportional for some positive r independent of n.

A *positive allocation* is an allocation with a positive egalitarian welfare - every agent receives a piece with a strictly positive value.

Given a social welfare function W and a fairness criterion F, the price-of-fairness relative to W and F (also called: "the W-price-of-F") is the ratio:

$$\frac{\max_X W(X)}{\max_{Y\in F} W(Y)} \qquad (*)$$

where the maximum at the nominator is over all allocations X and the maximum at the denominator is over all allocations Y that also satisfy the fairness criterion F. We are mainly interested here in the utilitarian-price-of-proportionality, utilitarian-price-of-partial-proportionality and the utilitarian-price-of-positivity.

When there are geometric constraints, such as connectivity, both maxima are taken only on allocations that satisfy these constraints. The price-of-fairness may be either higher or lower than with arbitrary pieces, so both upper bounds and lower bounds have to be re-calculated.

2 Related Work

2.1 Price of fairness

The trade-off between fairness and efficiency has been studied in various contexts, from allocation of computer network resources to allocation of landing-times to airplanes. Bertsimas et al [3] present several such problems and prove two generic bounds on two prices-of-fairness:

- The *utilitarian-price-of-proportional-fairness* is $\Theta(\sqrt{n})$ and this is tight. By "proportional fairness" they refer to a generalization of the Nash bargaining solution from two to n players. In the context of cake-cutting, a "proportionally fair" allocation is also proportional to the usual sense (each agent receives at least $1/n$).
- The *utilitarian-price-of-max-min-fairness* is $\Theta(n)$ and this is tight. By "max-min fairness" they refer to maximizing the egalitarian welfare.

Their bounds are valid whenever the space of utility vectors of all feasible allocations is compact and convex. By the Dubins-Spanier theorems [9], the space of utility vectors in cake-cutting with arbitrary (disconnected) pieces is indeed compact and convex. Hence, the bounds of [3] are valid in this case, too.

Caragiannis et al [6] study several different fair-allocation settings. In the cake-cutting setting, they prove the same bounds as [3] in different ways. In addition, they study fair allocation of indivisible items, as well as divisible and indivisible *chores* (when each agent wants to receive as little as possible).

Aumann et al [2] study the price-of-fairness in cake-cutting with connected pieces. In this setting, the space of utility vectors is not convex so the generic bounds of [3] do not apply. They prove, among other results, that the utilitarian price of proportionality is still $\Theta(\sqrt{n})$, as in the disconnected case. They also consider the efficiency-fairness trade-off from the other direction, and show that the *proportional-price-of-utilitarianism* might be infinite.

Zivan [20] studies a compromise between efficiency and fairness in the context of cake-cutting between two agents with disconnected pieces. He shows that it is possible to give both agents at least $(1-l)/2$ of their total cake value in an allocation which he calls "l-trust-efficient".

Finally, Arzi [1] shows that partially-proportional cake allocations can be much more efficient than proportional ones. In particular, she shows that in some cases, an $(1-\epsilon)$-proportional allocation, for any $\epsilon > 0$, can have a utilitarian welfare $\Theta(\sqrt{n})$ times greater than the maximum possible in a proportional allocation.

2.2 Partial proportionality

While proportionality is the most common criterion of fairness in cake-cutting, it is often relaxed to partial-proportionality in order to achieve additional goals besides improving the utilitarian welfare, such as:

1. Speed. Finding a proportional division takes $\Theta(n \log n)$ queries, but finding a $(1/T)$-proportional division takes only $\Theta(n)$ queries, for some sufficiently large $T \geq 10$ [10, 11].
2. Guaranteeing a minimum-size constraint. [7] prove that it is impossible to guarantee a $(1/T)$-proportional allocation for any finite T, and provide additive approximation algorithms.
3. Satisfying geometric constraints like square pieces [18, 17]. For example, when the cake is square and the pieces must be square, it is impossible to guarantee a $(1/T)$-proportional allocation for any $T \leq 2$, but there is an algorithm that guarantees a 1/4-proportional allocation.

3 Disconnected Pieces

In this section we assume that the agents may receive arbitrary disconnected pieces. Our main lemma is:

Lemma 1. *Given cake-allocations X and Y and a constant $r \in [0, 1]$, there exists an allocation Z such that, for every agent i:*

$$V_i(Z_i) \geq rV_i(Y_i) + (1 - r)V_i(X_i)$$

Proof. The Dubins-Spanier convexity theorem [9] says that the space of utilities of cake-allocations is convex. Hence, there exists an allocation Z such that $\forall i$: $V_i(Z_i) = rV_i(Y_i) + (1 - r)V_i(X_i)$.

Since the Dubins-Spanier theorem is not constructive, we give here a constructive protocol for creating the allocation Z when the ratio r is a rational number. Suppose $r = p/q$, where p, q are positive integers and $p < q$. For every pair of agents i, j, the protocol works as follows:

- Agent i divides the piece $X_i \cap Y_j$ to q pieces that are equal in its eyes.
- Agent j takes the p pieces that are best in its eyes.
- Agent i takes the remaining $q - p$ pieces.

Note that each for agent i, Z_i contains $X_i \cap Y_i$. All in all, Z_i contains nq pieces: np pieces that agent i took from other agents (including itself) in piece Y_i and $n(q - p)$ pieces that were left for agent i from other agents in piece X_i.

From every piece $Y_i \cap X_j$ (for $j \in \{1, \dots, n\}$), agent i picks the best p out of q pieces, which give it a value of at least $\frac{p}{q}V_i(Y_i \cap X_j)$. Hence, its total value from these np pieces is at least $rV_i(Y_i)$.

In addition, from every piece $X_i \cap Y_j$ (for $j \in \{1, \dots, n\}$), agent i receives $q-p$ out of q equal pieces, which give it a value of exactly $\frac{q-p}{q}V_i(X_i \cap Y_j)$. Hence, its total value from these $n(q - p)$ pieces is exactly $(1 - r)V_i(X_i)$. □

In the above proof, the utilitarian/egalitarian welfare of the allocation Z is at least r times the utilitarian/egalitarian welfare of Y and at least $1 - r$ times the utilitarian/egalitarian welfare of X. Suppose Y is a utilitarian allocation and X is a proportional allocation. Then, Z is simultaneously r-utilitarian and $(1 - r)$-proportional. Hence our Theorem 1 is proved and we get:

Corollary 1. *For every $r \in [0,1]$, the utilitarian price of $(1-r)$-proportionality is at most $1/r$.*

Figure 1 represents the tradeoff proved by Corollary 1. This is an upper bound on the tradeoff, i.e, in specific cases it may be possible to attain a smaller price-of-fairness. As expected, the curve approaches 1 at infinity, since the utilitarian price of proportionality is not bounded by any constant; Example 1 shows that it goes to infinity with the number of agents. The proportionality price of full utilitarianism is similarly unbounded. However, the curve shows that there are many options besides full utilitarianism and full proportionality. For example, the point (2,2), representing a 1/2-proportional 1/2-utilitarian allocation, seems a reasonable compromise.

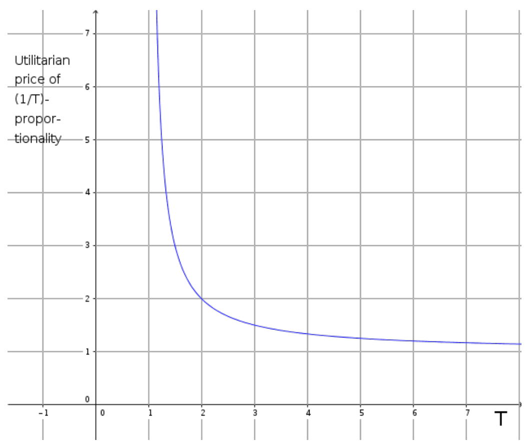

Fig. 1. A curve showing an upper bound on tradeoff between proportionality and utilitarian welfare. The height of the curve at $x = T$ is the utilitarian price of (1/T)-proportionality.

4 Connected Pieces - Negative Result

In this section we assume that each agent can derive utility only from a single connected interval. Since the protocol of Lemma 1 may give each agent as many as nq disjoint pieces (where q is the denominator of the parameter r), it cannot be used here. Moreover, Dubins-Spanier Convexity theorem (and Lemma 1) are no longer true - the space of utilities of cake-allocations is no longer convex. This is a corollary of the following counter-lemma:

Lemma 2. *Let X be a connected allocation. Then there exist n agents such that, in every positive connected allocation Z, for some agent i, $V_i(Z_i) \leq V_i(X_i)/n$.*

This lemma implies that there is no partially-proportional allocation in which all agents keep a constant fraction of the utility they had in X.

Before proving Lemma 2 we introduce some terminology.

Definition 1. *Given an agent i, an interval X_i and an integer l, we say that i has an l-**fractioned value-density** in X_i if there exists a set of l disjoint sub-intervals of X_i, the first and last of which touch the two endpoints of X_i, such that:*

- *Agent i values each interval as exactly $V_i(X_i)/l$;*
- *The value-density of i is zero in the $l-1$ "holes" between the intervals.*

The following illustration shows an l-fractioned value-density function on the interval $X_i = [0, 9]$, where $l = 5$:

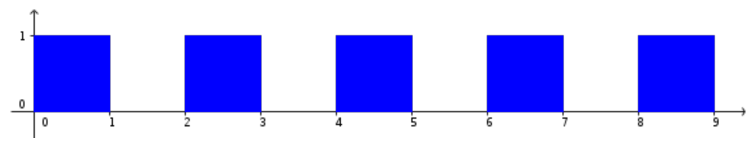

Note 1. If agent i has an l-fractioned value-density in X_i, and $V_i(Z_i) > V_i(X_i)/l$, then Z_i entirely contains one of the $l-1$ holes.

Proof (of Lemma 2). Without loss of generality we take $i = n$ and recall that X_n (an interval) is the share of agent n in the fixed allocation X. Suppose agent n has an n-fractioned value-density in X_n. Moreover, suppose the valuations of the other agents are such that each agent $j \in \{1, \dots, n-1\}$ has all its value concentrated in the j-th "hole" in X_n. By Note 1, if agent n receives value more than $V_n(X_n)/n$, then one of the other agents receives value 0. Hence, in any positive allocation Z, agent n must receive at most $V_n(X_n)/n$. □

In contrast to the disconnected case, here we cannot always convert an allocation X to a partially-proportional allocation in which *each agent* receives at least a constant fraction of its utility in X.

Note that the allocation used in the proof of Lemma 2 is not utilitarian-optimal. In particular, instead of giving X_n to agent n and getting a utilitarian welfare of 1, we can divide it among the other $n-1$ agents and get a utilitarian welfare of $n-1$. Below we show a more complicated construction that gives an analogous result for utilitarian-optimal allocations.

Lemma 3. *Let X be a **utilitarian** connected allocation. There exists a cake-cutting instance in which, in every positive connected allocation Z, for some agent i, $V_i(Z_i) \leq V_i(X_i) \cdot \Theta(\lg \lg n / \lg n)$.*

Proof. Let k be an integer such that $n > k + k^k$. Let $m = k - 1$ (This implies $k, m \in \Theta(\lg / \lg n \lg n)$). Consider the following cake-cutting instance.

(1) There are k *rich agents*. For each rich agent i, $V_i(X_i) = V_i(C) = 1$, i.e, all its value-density is concentrated in X_i (its share in allocation X). Its value

density there is m-fractioned (so that each fraction is worth $1/m$). Note that in each region X_i there are $m-1$ "holes" in which the value-density of i is 0. The total number of holes is $k(m-1)$.

(2) There are $(m-1)^k$ *poor agents*. Each poor agents has all its value-density scattered among k holes (so that each hole is worth $1/k$). The value-density of a poor agent is uniform within each hole. Note that for each hole there may be several poor agents who want that hole.

We now prove that X is utilitarian. Since $m < k$, $1/m > 1/k$. Giving a hole to one or more poor agents gives these agents a total utility of at most $1/k$, but subtracts at least $1/m$ from a rich agent. Hence, in a utilitarian allocation, poor agents receive nothing.

Let Z be a positive allocation. Our goal is to prove that in Z, at least one rich agent receives a value of at most $1/m$. Suppose as contradiction that each rich agent receives more than $1/m$. Then, by Note 1, in each region at least one hole is covered. All in all, k holes are covered. The number of holes in each region is $m-1$, so the number of subsets of k holes is $(m-1)^k$, which equals the number of poor agents. Assume that each poor agent wants a different subset of k holes. Then, at least one poor agent remains with no value. Hence, if Z is positive then some rich agent receives a value of at most $1/m$, which is $V_i(X_i) \cdot \Theta(\lg \lg n / \lg n)$. □

We use a similar construction to prove a lower bound on the utilitarian price of giving each agent a positive value.

Lemma 4. *The price of positivity may be as high as $\Theta(\sqrt{\frac{\lg n}{\lg \lg n}})$.*

Proof. Let k be an integer such that $n > k + \binom{k(m-1)}{k}$. Let $m = \sqrt{k}$ (this implies $m \in \Theta(\sqrt{\lg n / \lg \lg n})$). Consider the following cake-cutting instance.

There are k regions and k *rich agents*, each of whom has an m-fractioned value-density in its region.

There are $\binom{k(m-1)}{k}$ *poor agents*, each of whom has its value-density scattered among k holes.

Consider now any positive allocation Z. Our goal is to prove that in Z, the utilitarian welfare is at most $O(1/\sqrt{k})$ of the optimum. Suppose as contradiction that all rich agents together receive a value of more than $(2k-1)/m$. This means that they receive $2k-1$ fractions in k regions, so in all regions together at least k holes are covered. The total number of holes is $k(m-1)$ and the total number of different hole subsets is $\binom{k(m-1)}{k}$ which is the number of poor agents. If each poor agent wants a different subset of k holes, then at least one poor agent remains with no value. Since Z is positive, the sum of utilities of all rich agents in Z must be at most $(2k-1)/m$.

Moreover, the sum of utilities of all poor agents in each hole is $1/k$, in each region - less than m/k, and overall - less than m. Hence, the total utilitarian welfare in a partially-proportional division is less than $m + (2k-1)/m$. But the utilitarian welfare in the utilitarian allocation X is k (each rich agent receives

1). Hence, the utilitarian-welfare approximation ratio is at most:

$$\frac{m + (2k-1)/m}{k} < \frac{m^2 + 2k}{mk} = \frac{m}{k} + \frac{2}{m} = \Theta(1/\sqrt{k}) = \Theta(\sqrt{\lg\lg n/\lg n})$$

as claimed. □

Lemma 4 also proves our Theorem 2.

5 Connected Pieces - Positive Result

We start with a positive counterpart of Lemma 3.

Lemma 5. *Given a connected utilitarian allocation X to n agents, there exists a connected allocation Z such that, for every agent i, the following guarantees hold simultaneously:*

$$V_i(Z_i) \geq V_i(X_i)/O(\log n)$$
$$V_i(Z_i) > 0$$

Proof. Let $m = \lceil \log_2 n \rceil + 1$. Divide the agents to two groups:

- *Rich agents* - whose value in X is positive;
- *Poor agents* - whose value in X is zero.

Let k be the number of rich agents; number them $1, \ldots, k$.

Ask every rich agent $i \in \{1, \ldots, k\}$ to divide its plot X_i to m intervals with equal value. This results in mk intervals. The value of each interval in X_i for agent i is $V_i(X_i)/m$, which is at most $1/m$.

For each rich agent i, define the *extreme intervals* as the leftmost and the rightmost interval in X_i. the total number of extreme intervals is $2k$. The value of each extreme interval to its owner is at most $1/m$. Because X is utilitarian, the value of each extreme interval to any other agent j is at most $1/m + V_j(X_j)$ - otherwise we could improve the utilitarian welfare by taking this interval from i and giving it to j. In particular, any poor agent values any extreme interval as at most $1/m$.

We want to create an allocation Z in which each rich agent receives one of its two extreme intervals; this will guarantee each rich agent a value of at least $V_i(X_i)/m$. The total number of such selections is 2^k. Out of these selections, we would like to select one in which the remainder has a positive value for each poor agent. We call such a selection "good". Conversely, a selection in which the remainder has zero value for some poor agent is called "bad". Our goal is to prove that there is at least one good selection.

There are several easy cases:

Easy case #1: $k \leq m - 1$. Then, in any selection of extreme-interval-per-rich-agent, the total value that is made unavailable to the zero agents is at most $k \cdot (1/m) \leq 1 - 1/m$, so the remaining cake (after the allocation to the rich agents) has a value of at least $1/m$ to each zero agent. All selections are good.

Easy case #2: for some poor agent j and some rich agent i, both extreme intervals in X_i have positive V_j. Then, in any selection of extreme-interval-per-rich-agent, the remainder will have positive V_j.

Easy case #3: for some poor agent j and some rich agent i, the part of X_i not contained in the extreme intervals have positive V_j. Then, in any selection of extreme-interval-per-rich-agent, the remainder will have positive V_j.

Therefore, from now on we assume that none of these easy cases happen. That is, we assume: (1) $k \geq m$, (2) for each poor agent j and rich agent i, one extreme-interval of X_i has $V_j = 0$, and the other one has $V_j \leq 1/m$.

Consider some poor agent j, and let l be the number of extreme-intervals in which V_j is positive. By the previous paragraph, $m \leq l \leq k$. The total number of ways to select an extreme interval per rich agent is 2^k. Out of these, the total number of selections in which all of j's positive intervals are taken, is at most $2^{k-l} \leq 2^{k-m}$. The total number of poor agents is at most n. By the union bound, the total number of selections in which all positive intervals of some poor agent are taken (we called such selections "bad"), is at most $n \cdot 2^{k-m} = 2^k \cdot (n/2^m)$. By definition of m, $n < 2^m$, hence the total number of bad selections is less than 2^k, hence there exists at least one good selection.

The number of steps required to find a good selection is exponential in k, but this is not relevant since we are only interested in existence result. After implementing a good selection, the remainder contains at most $k + 1$ intervals and it has a positive value for each poor agent. It can be divided among the poor agents using any algorithm for dividing multiple cakes, e.g. the "archipelago division algorithm" [18]. In the resulting division, each agent has value at least $V_i(X_i)/m$ and all agents have positive value.

Corollary 2. *The utilitarian price of positivity is at most* $1/O(\log n)$.

This also proves our Theorem 3.

5.1 Future Work

We presented a quite complete efficiency-fairness tradeoff curve for cake-cutting with disconnected pieces. The main challenge for future work is to complete the efficiency-fairness tradeoff curve for cake-cutting with connected pieces. In particular, it is interesting to know what reduction in utilitarian welfare is sufficient for guaranteeing: (1) a positive value for each agent, (2) a constant fraction of $1/n$ for each agent?

An additional possible generalization is to allow each agent to receive a constant number of intervals. This seems like a fair compromise between arbitrarily-many intervals and a single interval. As far as I know, it has not been studied.

References

1. Arzi, O.: Cake Cutting: Achieving Efficiency While Maintaining Fairness. Master's thesis, Bar-Ilan University (Oct 2012), under the supervision of Prof. Yonatan Aumann

2. Aumann, Y., Dombb, Y.: The Efficiency of Fair Division with Connected Pieces. Web, Internet and Network Economics 6484, 26–37 (2010), http://ccc.cs.uni-duesseldorf.de/COMSOC-2010/papers/comsoc2010.pdf#page=231
3. Bertsimas, D., Farias, V.F., Trichakis, N.: The Price of Fairness. Operations Research 59(1), 17–31 (Feb 2011), http://pubsonline.informs.org/doi/abs/10.1287/opre.1100.0865
4. Brams, S.J.: Mathematics and Democracy: Designing Better Voting and Fair-Division Procedures. Princeton University Press, first edition edn. (Dec 2007), http://www.sciencedirect.com/science/article/pii/S089571770800174X
5. Brams, S.J., Taylor, A.D.: Fair Division: From Cake-Cutting to Dispute Resolution. Cambridge University Press (Feb 1996), http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0521556449
6. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M.: The Efficiency of Fair Division. Theory of Computing Systems 50(4), 589–610 (Sep 2012), http://dx.doi.org/10.1007/s00224-011-9359-y
7. Caragiannis, I., Lai, J.K., Procaccia, A.D.: Towards more expressive cake cutting. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11). pp. 127–132. AAAI Press (2011), http://www.cs.cmu.edu/˜./arielpro/papers/nonadd.ijcai11.pdf
8. Chevaleyre, Y., Dunne, P.E., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodriguez-Aguilar, J.A., Sousa, P.: Issues in Multiagent Resource Allocation. Informatica 30(1), 3–31 (2006), http://www.informatica.si/index.php/informatica/article/view/70
9. Dubins, L.E., Spanier, E.H.: How to Cut A Cake Fairly. The American Mathematical Monthly 68(1), 1+ (Jan 1961), http://dx.doi.org/10.2307/2311357
10. Edmonds, J., Pruhs, K.: Balanced Allocations of Cake. In: FOCS. vol. 47, pp. 623–634. IEEE Computer Society (Oct 2006), http://dx.doi.org/10.1109/focs.2006.17
11. Edmonds, J., Pruhs, K., Solanki, J.: Confidently Cutting a Cake into Approximately Fair Pieces. Algorithmic Aspects in Information and Management pp. 155–164 (2008), http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.145.8396&rep=rep1&type=pdf
12. Mill, J.S.: Utilitarianism. Hackett Pub Co, 2 edn. (1863), http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/087220605X
13. Moulin, H.: Fair Division and Collective Welfare. The MIT Press (Aug 2004), http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0262633116
14. Procaccia, A.D.: Cake Cutting Algorithms. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.) Handbook of Computational Social Choice, chap. 13. Cambridge University Press (2015), http://www.cs.cmu.edu/˜arielpro/papers/cakechapter.pdf
15. Rawls, J.: A Theory of Justice. Belknap Press of Harvard University Press, revised edition edn. (1971), http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0674000781
16. Robertson, J.M., Webb, W.A.: Cake-Cutting Algorithms: Be Fair if You Can. A K Peters/CRC Press, first edn. (Jul 1998), http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/1568810768

17. Segal-Halevi, E., Hassidim, A., Aumann, Y.: Envy-Free Cake-Cutting in Two Dimensions. In: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-15). pp. 1021–1028 (Jan 2015), http://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9656
18. Segal-Halevi, E., Nitzan, S., Hassidim, A., Aumann, Y.: Fair and Square: Cake-Cutting in Two Dimensions (Oct 2015), http://arxiv.org/abs/1510.03170, arXiv preprint 1510.03170
19. Steinhaus, H.: The problem of fair division. Econometrica 16(1), 101–104 (Jan 1948), http://www.jstor.org/stable/1914289
20. Zivan, R.: Can trust increase the efficiency of cake cutting algorithms. In: The International Conference on Autonomous Agents and Multiagent Systems (AAMAS). pp. 1145–1146 (2011), http://aamas.csc.liv.ac.uk/Proceedings/aamas2011/papers/B2.pdf