

Envy-Free Cake-Cutting in Two Dimensions

Erel Segal-Halevi, Avinatan Hassidim, Yonatan Aumann

Bar-Ilan University
Ramat-Gan
Israel 5290002

Abstract

We consider the problem of fair division of a two dimensional heterogeneous good among several agents. Applications include division of land as well as ad space in print and electronic media. Classical cake cutting protocols either consider a one-dimensional resource, or allocate each agent several disconnected pieces. In practice, however, the two dimensional shape of the allotted piece is of crucial importance in many applications, e.g., squares or bounded aspect-ratio rectangles are most useful for building houses as well as advertisements. We thus introduce and study the problem of envy-free two-dimensional division wherein the utility of the agents depends on the geometric shape of the allocated pieces (as well as the location and size). In addition to envy-freeness, we require that the fraction allocated to each agent be at least a certain constant that depends only on the shape of the cake and the number of agents. We focus on the case where the allotted pieces must be square and the cakes are either squares or the unbounded plane. We provide algorithms for the problem for settings with two and three agents.

1 Introduction

Fair division (often termed *cake-cutting*) is an active field of research and application in mathematics, economics, and recently also in AI. The basic setting considers a heterogeneous good, e.g. land, that must be divided among several agents. The agents may have different preferences over the possible pieces/sub-sets of the good, e.g. one agent prefers the forest while the other prefers the sea shore, and the goal is to divide the good among the agents in a way that deemed “fair”. *Fairness* can be defined in several ways, but, *proportionality* and *envy-freeness* are the most commonly used. Proportionality means that each agent gets at least its “fair-share” of the good, i.e. with n agents, the piece allotted to each agent is worth at least $1/n$ of the value of the entire good - according to agent’s subjective valuations. Envy-freeness means that no agent would prefer getting a piece allotted to another agent. The existence of proportional divisions was already proved in the initial work of (Steinhaus 1948), and the existence of envy-free divisions was established by (Stromquist 1980). This latter proof is existential in nature. *Constructing* such envy-free divisions turns out to

be a much harder task. The first protocol for three agents, involving ten steps and five cuts, was discovered by Selfridge and Conway (Brams and Taylor 1996). The construction of a protocol for envy-free division among four or more agents was a long-standing open problem, resolved only in 1995 with the publication of the Brams-Taylor protocol (Brams and Taylor 1995).

Interestingly, in all this research, spanning over half a century, little attention has been given to the case of dividing a *two-dimensional* good.

Indeed, almost all work on cake cutting explicitly assumes that the cake is a one-dimensional interval, and the allotted pieces are either sub-intervals, or a (possibly unbounded) collection thereof. This is usually justified by the reasoning that higher dimensional settings can always be projected onto one dimension, and hence, fairness in one-dimension implies fairness in higher dimensions. However, projecting back from the one dimension, the resulting two-dimensional plots are thin rectangular slivers, of little use in most practical applications; it is hard to build a house on a $10 \times 1,000$ meter plot even though its area is a full hectare, and a thin 0.1-inch wide advertisement space would ill-serve most advertisers, regardless of its height. Thus, in most applications, the geometric *shape* of the allotted piece is of prime importance.

Hence, we argue that when dividing a two-dimensional resource, one must also require that allotted plots be of “usable” shape, e.g. a square. Equivalently, one can allow arbitrary plots but note that the utility of any such plot is determined by the most valuable square therein. These two formulations are equivalent and we use them interchangeably throughout the paper.

Once allotted plots must be of a specific shape(s), it may no longer be possible to allocate the entire land; a square cannot be divided into two squares. Thus, we must assume *free disposal*, meaning that the protocol need not allocate the entire “cake”. Once free disposal is permitted, however, envy-freeness can trivially be obtained by simply giving nothing to all agents, which is clearly not a solution of interest. Thus, our goal is to obtain *an envy-free division that also guarantees some minimum utility to each agent*. Specifically, we seek an envy-free division that guarantees that each agent’s piece is worth at least an α -fraction of the value of the entire cake, according to the agents subjective

valuations. In general, we want α as large as possible, but note that apriori it is not at all clear that such a guarantee is possible for any constant α .

We focus on the cases of two and three agents (which, unsurprisingly, already turn out to be non-trivial). We focus on the case where the usable shapes are square, and the initial land is either also a square or the unbounded plane (when the cake is unbounded we also allow the squares to be unbounded, i.e. we consider a quarter-plane and a half-plane as 'squares' of infinite side-length). Similar results, though more complex in presentation, can be obtained for other shapes.

1.1 Motivating example

The following example illustrates the insufficiency of existing cake-cutting algorithms. Suppose you and a partner are joint owners of a square land estate near the sea. The estate is a 100-by-100 square meters whose western side is adjacent to the sea. Suppose you believe that the most valuable part of the estate is the seashore, e.g., because your desire is to have a home near the sea. One day you decide to break the partnership and divide the land using the classic procedure for envy-free division: "You cut, I choose". You let your partner divide the land to two plots, knowing that you have the right to choose the plot that is more valuable according to your own subjective preferences. Your partner makes a cut parallel to the shoreline at a distance of only 1 meter from the sea.¹ Which of the two plots would you choose? The western plot contains a lot of sea shore, but it is so narrow that it has no room for building anything. On the other hand, the eastern plot is large but does not contain any shore land. Whichever plot you choose, the division will not be proportional for you, because your utility is far less than half the value of the original land estate. Of course the cake *could* be cut in a more sensible way (e.g. by a line perpendicular to the sea), but the current protocols say nothing about how exactly the cake should be cut in each situation in order to guarantee that the division is fair in a way that respects the geometric constraints.

The reason that the classic procedure fails here is that it is based on the assumption of *additive utilities*, meaning that the sum of the utilities you can derive from two parts of the land equals the utility you derive from the entire land. If utilities were additive then one of the plots would necessarily have a utility of at least $\frac{1}{2}$ of the total value. However, in reality the utility of land is not an additive function.

In the above example, assuming you want to build a square house, the utility you can derive from each land-plot is the utility of the most valuable square contained in that plot, and thus the sum of utilities of the western and eastern parts in separation is much less than the utility of the land-estate as a whole. Thus, while the cut-and-choose protocol guarantees an envy-free division, it does not guarantee any positive utility to the agents.

¹The reason why he decided to cut this way is irrelevant since a fair division protocol is expected to guarantee that the division is fair for every agent playing by the rules, regardless of what the other agents do.

1.2 Our results

For the case that the original cake is a square, we present:

- A protocol for envy free division with two agents that guarantees each agent a utility of at least $1/4$ -th of the value of the entire cake. This matches the existential upper bound established by (Segal-Halevi, Hassidim, and Aumann 2014).
- A protocol for envy free division with three agents that guarantees each agent a utility of at least $1/10$ -th of the value of the entire cake. The existential upper bound of (Segal-Halevi, Hassidim, and Aumann 2014) for this case is $1/6$ -th.

For the case that the original cake is the unbounded plane, we present a protocol for envy free division with three agents that guarantees each agent a utility of at least $1/3$ -rd of the value of the entire cake.² This protocol is clearly existentially optimal, as $1/n$ is the best possible fraction in the case where all agents have identical valuations.

The remainder of the paper is structured as follows. Immediately following we review the related research. The formal definitions and model are provided in Section 2. Section 3 presents the core geometric concepts and techniques, including our main concept of a *knife function*, which is a generalization of the "moving knife" concept used by both classic and modern works on cake cutting (Dubins and Spanier 1961; Stromquist 1980; Brams, Taylor, and Zwicker 1997; Manabe and Okamoto 2012). These geometric techniques are then applied in the construction of the envy-free division procedures for two agents (Section 4) and three agents (Section 5).

1.3 Related work

As far as we know, existing protocols for envy-free division (Stromquist 1980; Brams and Taylor 1995; Reijnierse and Potters 1998; Su 1999; Barbanel and Brams 2004; Manabe and Okamoto 2010; Cohler et al. 2011; Deng, Qi, and Saberi 2012; Kurokawa, Lai, and Procaccia 2013; Chen et al. 2013) do not make any shape-related guarantees. When applied to agents with non-additive utility functions, the utility per agent might be arbitrarily small.

Most other results related to cake-cutting with non-additive utilities is either purely existential (Sagara and Vlach 2005; Dall'Aglio and Maccheroni 2009; Hüsseinov and Sagara 2013) or assume a 1-dimensional cake (Su 1999; Caragiannis, Lai, and Procaccia 2011; Mirchandani 2013)

Relatively few papers explicitly relate to a two-dimensional cake.³ Two of them discuss the problem of dividing a disputed territory between several bordering countries, with the constraint that each country should get a piece that is adjacent to its border: (Hill 1983) proved that such a partition exists and (Beck 1987) provided a division procedure. (Iyer and Huhns 2009) describe a procedure that asks

²For two agents, the classical cut-and-choose procedure works.

³Several authors studied a circular cake (Thomson 2007; Brams, Jones, and Klamler 2008; Barbanel, Brams, and Stromquist 2009), but this is a one-dimensional circle and the pieces are one-dimensional arcs corresponding to thin wedge-like slivers.

each of the n agents to draw n disjoint rectangles on the map of the two-dimensional cake. These rectangles are supposed to represent the “desired areas” of the agent. The procedure tries to give each agent one of his n desired areas. However, it does not succeed unless each rectangle proposed by an individual intersects at most one other rectangle drawn by any other agent. If even a single rectangle of Alice intersects two rectangles of George, then the procedure fails and no agent gets any piece.

2 The Model

A cake C , which is a measurable subset of the two-dimensional Euclidean plane, has to be divided among n agents. Every agent $i \in \{1, \dots, n\}$ should receive a measurable piece $P_i \subseteq C$ (Note that we do not require that the entire cake is divided).

Every agent i has an (additive) value measure V_i over parts of C , normalized such that $V_i(C) = 1$. The value measures are assumed to be finite and non-negative. The measures are also assumed to be *absolutely continuous with respect to area* (or just *continuous* for short), i.e., any piece with an area of 0 has a value of 0 (Hill and Morrison 2010). Hence, the value of a plot is the same whether or not it contains its boundary. Equivalently, for every $\epsilon > 0$ there is a $\delta > 0$ such that, for every s having area less than δ , the value $V(s)$ is less than ϵ .

There is a pre-defined family of shapes S , which is the family of *usable geometric shapes*. In this paper, S is the family of squares. Based on V_i and S we define the following shape-based utility function, which assigns to a piece $P \subseteq C$ the value of the most valuable square contained in P :

$$V^S(P) = \sup_{s \in S \text{ and } s \subseteq P} V(s)$$

For example, if Alice wants to build a square house but gets a land-plot P which is not square, then her utility is determined by the most valuable square contained in her plot P . From now on, we use the term *value* to refer to the additive measure V and the term *utility* to refer to the (not necessarily additive) function V^S .

An *envy-free partition* of a cake C is a partition in which the utility of an agent from his allocated piece is at least as large as his utility from every other:

$$\forall i, j \in \{1, \dots, n\} : V_i^S(P_i) \geq V_i^S(P_j)$$

In addition to envy-freeness, every partition can be characterized by its level of *proportionality*, which is the utility of the least fortunate agent (also known as *egalitarian social welfare*):

$$Prop(C, S, \{V_i\}_{i=1}^n, \{P_i\}_{i=1}^n) = \min_{i \in \{1, \dots, n\}} V_i^S(P_i)$$

So a proportional partition is a partition with a proportionality of at least $\frac{1}{n}$.

We are interested in finding, for given C , n and S , the largest proportionality level that can be attained for every combination of continuous value measures in an envy-free




Figure 1: Geometric loss factors relative to the family of squares.

division. We call this number the *n-agent envy-free proportionality level* of C and S :

$$Prop_E(C, n, S) = \inf_{V_i} \sup_{P_i} Prop(C, S, \{V_i\}_{i=1}^n, \{P_i\}_{i=1}^n)$$

where the infimum is taken over all n -tuples of continuous value measures V_i and the supremum is taken over *all envy-free partitions* of C to n agents.

A similar function, $Prop(C, n, S)$, can be defined exactly the same as $Prop_E(C, n, S)$ with the only difference being that the supremum is taken over *all partitions* of C (regardless of envy). Obviously, because the supremum in $Prop_E$ is taken over a smaller set:

$$Prop_E(C, n, S) \leq Prop(C, n, S)$$

Applying this notation, classic cake-cutting results imply that for every cake C :

$$Prop(C, n, All) = Prop_E(C, n, All) = \frac{1}{n}$$

Where “All” is the collection of all geometric shapes. That is: when all geometric shapes are usable, every cake C can be divided in an envy-free way such that each agent receives $\frac{1}{n}$ of the total utility, for every combination of continuous value measures. Recently, (Segal-Halevi, Hassidim, and Aumann 2014) proved that:

$$\frac{1}{4n-4} \leq Prop(Square, n, Squares) \leq \frac{1}{2n}$$

Our contribution in this paper is to calculate a lower bound on $Prop_E(Square, n, Squares)$ and $Prop_E(Plane, n, Squares)$ for $n \in \{2, 3\}$. Note that envy-free division is a much more difficult task than proportional division even without geometric constraints, and even when there are only 3 agents.

3 Geometric Concepts

The land-estate example described in the introduction hints that to achieve a fair division we must constrain the ways in which agents are allowed to cut the cake. To this end we now define several properties of geometric shapes.

3.1 Geometric loss

Definition 1. For a cake C and family of shapes S , the *geometric loss* factor of C relative to S is the maximum factor by which the utility of an agent from C is reduced by his insistence on using shapes only from family S . Formally:

$$Loss(C, S) = \sup_V \frac{V(C)}{V^S(C)}$$

where the supremum is over all continuous finite measures V having $V(C) > 0$. The minimum possible loss factor is 1 which means no loss. This is always the case when $C \in S$ since in that case $V^S(C) = V(C)$. When C is not in S , the loss is generally larger than 1. For example, assume that C is a 30-by-20 rectangle. The largest square contained in C is 20-by-20. Hence, if V is uniform over C (as in Figure 1/Left, representing an agent who wants a maximal amount of land), then $\frac{V(C)}{V^S(C)} = \frac{600}{400} = \frac{3}{2}$, implying that $Loss(C, Squares) \geq \frac{3}{2}$. But the loss may be larger: suppose V is uniform over the “shores” in the east and west sides of C (as in Figure 1/Right, representing an agent who wants a square with a maximal amount of shore land). In this case $\frac{V(C)}{V^S(C)} = 2$, implying that $Loss(C, Squares) \geq 2$. We will see later that the loss in this case is exactly 2, and in general the geometric loss of a rectangle with a length/width ratio of L is $\lceil L \rceil$, so a thinner rectangle has a larger geometric loss.

The importance of the geometric loss concept becomes clear when it is generalized to cake partitions:

Definition 2. For a partition of a cake to pieces $P_1 \sqcup P_2 \sqcup \dots \sqcup P_m = C$ and a family of shapes S , the *geometric loss of the partition* is the sum of the geometric loss factors of the pieces:

$$Loss(\{P_1, P_2, \dots, P_m\}, S) = \sum_{i=1}^m Loss(P_i, S)$$

Intuitively, lower geometric loss is better: when a cake is partitioned with a low geometric loss, it is possible to choose at least one piece with a sufficiently high utility. For example, assume the cake C is a 100-by-100 square and S is the family of squares. If C is partitioned near its boundary to 100-by-1 and 100-by-99 rectangles (like the land-estate in the introduction), then the loss factors of the pieces are 100 and 2 respectively so the loss of the partition is 102. But if C is partitioned in the middle to two 100-by-50 rectangles, then the loss factor of both pieces is 2 so the loss of the partition is 4, in accordance with our intuition that the second partition is “fairer” for the chooser. The following lemma formalizes this intuition.

3.2 Chooser lemma

Lemma. (Chooser Lemma) *If a cake C is partitioned to disjoint pieces $P_1 \sqcup \dots \sqcup P_m = C$ and the geometric loss of the partition is M , then for every value measure V and every family of shapes S :*

$$\max(V^S(P_1), \dots, V^S(P_m)) \geq \frac{V(C)}{M}$$




Figure 2: Cover numbers of several geometric shapes.

Proof. By additivity, $\sum_{i=1}^m V(P_i) = V(C)$. Multiplying the left-hand side by the geometric loss (M) and the right-hand side by the definition of geometric loss gives: $\sum_{i=1}^m V(P_i) \cdot M = \sum_{i=1}^m Loss(P_i, S) \cdot V(C)$. At least one of the m elements in the left-hand side must be greater than or equal to the corresponding element in the right-hand side. I.e., there is an i for which: $V(P_i) \cdot M \geq Loss(P_i, S) \cdot V(C)$. By Definition 1: $V^S(P_i) \geq \frac{V(P_i)}{Loss(P_i, S)} \geq \frac{V(C)}{M}$. \square

Hence, when the geometric loss of a partition is M , the chooser can choose a piece with a utility of at least $\frac{V(C)}{M}$.

3.3 Cover numbers

An upper bound on the geometric loss factor can be proved based on the following definition:

Definition 3. For a cake C and family of geometric shapes S , $CoverNum(C, S)$ is the minimum number of shapes from family S whose union is exactly C .

Some examples are depicted in Figure 2. When C is a hole-free rectilinear polygon and S the family of squares, a minimum cover can be found in polynomial time (Bar-Yehuda and Ben-Hanoch 1996).

Claim 1. For every cake C and family S :

$$Loss(C, S) \leq CoverNum(C, S)$$

Proof. Let $n = CoverNum(C, S)$ and let $P_1, \dots, P_n \in S$ be shapes that cover the cake C such that $C = P_1 \cup P_2 \cup \dots \cup P_n$. Let V be any value measure. A measure is additive, so $V(C) = V(P_1) + V(P_2) + \dots + V(P_n)$. Hence, there is at least one piece $P_i \in S$ such that $V(P_i) \geq \frac{V(C)}{n}$. Therefore, $V^S(C) = \sup_{p \subseteq C \text{ and } p \in S} V(p) \geq V(P_i) \geq \frac{V(C)}{n}$, which implies $Loss(C, S) \leq n$. \square

An immediate consequence of Definition 1 is that for every value measure V : $V^S(C) \geq \frac{V(C)}{Loss(C, S)} \geq \frac{1}{CoverNum(C, S)}$ (since $V(C)$ is normalized to 1). Going back to the 30×20 rectangle, as shown in Figure 2/Right, $CoverNum(C, Squares) = 2$ so $V^S(C) \geq \frac{1}{2}$, meaning that every agent, regardless of his preferences, can get a utility of at least half the total cake value.

3.4 S-Continuity

We are going to cut cakes using a generalization of a moving knife. To use this generalization correctly we have to make sure that the utility functions of the agents change continuously during the movement of the knife. To this end we present the following definition:

Definition 4. Let S be a family of shapes. A function $K(t)$ from a real interval to subsets of R^2 is called *S-continuous* if for every $\epsilon > 0$ there exists $\delta > 0$ such that, for every t and t' which are at most δ apart (i.e. $|t' - t| < \delta$), and for every shape $s' \in S$ contained in $K(t')$, there exists a shape $s \in S$ contained in $K(t)$ such that $s \subseteq s'$ and $Area(s' \setminus s) < \epsilon$.

Intuitively, *S-continuity* means that, as the function K changes over time, usable shapes are not created or destroyed “all of a sudden”, but rather grow or shrink in a smooth manner. As the examples below demonstrate, *S-continuity* is different from continuity.

(a) Let S be the family of squares parallel to the axes. The function $K_1(t) = [0, t] \times [0, 1]$, defined for $t \geq 0$, is *S-continuous*. Proof: Intuitively, the function $K_1(t)$ describes a rectangle growing smoothly eastwards; it is apparent that no squares with positive area are created all of a sudden. Formally, given $\epsilon > 0$, select a δ such that $2\delta + \delta^2 < \epsilon$. For every t, t' with $|t' - t| < \delta$, for every axis-parallel square in $K_1(t')$ with side-length $a + \delta$, there is a contained square in $K_1(t)$ with side-length at least a . The difference between these squares has an area of at most $2\delta + \delta^2 < \epsilon$.

(b) Similarly, the function $K_2(t) = [0, t] \times [0, t]$, describing a square growing from the origin, is *S-continuous*.

(c) In contrast, the function $K_3(t) = [0, t] \times [0, 1] \cup [1 - t, 1] \times [0, 1]$, defined for $t \in [0, 1]$, is not *S-continuous*. Proof: Intuitively, a square of side-length 1 is created at time $t = 0.5$, when the two components of $K_3(t)$ meet. Formally, let $\epsilon = 0.75$. For every $\delta > 0$, select $t = 0.5 - \frac{\delta}{3}$ and $t' = 0.5 + \frac{\delta}{3}$. Then $K_3(t')$ contains the square $s' = [0, 1] \times [0, 1]$, but all squares $s \subseteq K_3(t)$ have a side-length of less than 0.5, hence $Area(s' \setminus s) > 0.75 = \epsilon$. Note that K_3 is continuous in the following sense: the symmetric difference between $K_3(t + \delta)$ and $K_3(t)$ has zero area as δ tends to zero.

(d) Similarly, if S is the family of axis-parallel squares with side-length at least 0.5, then the function $K_1(t) = [0, t] \times [0, 1]$ is not *S-continuous*.

Lemma. (S-continuity lemma) *If the measure V is absolutely continuous and the function $K(t)$ is S-continuous, then the real function $v(t) := V^S(K(t))$ is uniformly continuous.*

Proof. Given $\epsilon > 0$, we have to show the existence of $\delta > 0$ such that, for every t, t' , if $|t' - t| < \delta$ then $|v(t') - v(t)| < \epsilon$.

By the absolute continuity of V , there is an $\epsilon' > 0$ such that every s having $Area(s) < \epsilon'$ has $V(s) < \epsilon$.

By the *S-continuity* of $K(t)$, there is a $\delta > 0$ such that, for every t, t' , if $|t' - t| < \delta$ then for every *S*-shape $s' \subseteq K(t')$, there is an *S*-shape $s \subseteq K(t)$ such that $s \subseteq s'$ and the difference $s' \setminus s$ has area less than ϵ' . Hence, $V(s') - V(s) = V(s' \setminus s) < \epsilon$.

By definition, $v(t') = V^S(K(t')) = V(s')$, where s' is an *S*-shape contained in $K(t')$ for which $V(s')$ is maximized.

By the same definition, $v(t) = V^S(K(t)) \geq V(s)$, where s is an *S*-shape contained in $K(t)$ having $V(s') - V(s) < \epsilon$. Hence, $v(t) > v(t') - \epsilon$. By a symmetric argument, $v(t') > v(t) - \epsilon$. \square

3.5 Knife functions

Definition 5. Let C be a cake. A *knife function* on C is a function from $[0, 1]$ to subsets of C having the following properties:

1. $K_C(0) = \emptyset$ and $K_C(1) = C$;
2. K_C is monotonically increasing with t , i.e. for every $t' > t$: $K_C(t') \supset K_C(t)$;
3. Both $K_C(t)$ and $C \setminus K_C(t)$ are *S-continuous* functions (see definition 4).

The name “knife function” comes from the classic moving knife procedure for proportional cake-cutting. When a knife moves over a cake, in the manner described by (Dubins and Spanier 1961), starting at time 0 and ending at time 1, the part of the cake already covered by the knife is a knife function according to our definition. However, our definition is more general and allows a “knife” that is not a straight line and does not necessarily move parallel to itself.

For example, let C be the unit square and S the family of squares. consider the two *S-continuous* functions defined in the previous section: $K_1(t) = [0, t] \times [0, 1]$ and $K_2(t) = [0, t] \times [0, t]$. It is easy to check that both of them are knife functions.

3.6 Geometric loss of knife functions

When a knife function K_C , defined on a cake C , is “stopped” at a certain time $t \in [0, 1]$, it induces a partition of C to the part which was already covered by the knife, $K_C(t)$, and the part not covered, $C \setminus K_C(t)$. Based on this partition, we can define the geometric loss of the knife:


Definition 6. Let C be a cake. For every knife function K_C , define its *geometric loss relative to the family S* as:

$$Loss(K_C, S) = \max_t Loss(K_C(t), S) + Loss(C \setminus K_C(t), S)$$


Intuitively, if $Loss(K_C, S) = M$ then the knife can be stopped at any time t and the resulting partition has a geometric loss of at most M . Two examples are illustrated in Figure 3, from left to right:

(a) Let $C = [0, L] \times [0, 1]$. Define the following knife function: $K_C(t) = [0, L] \times [0, t]$. Obviously both $K_C(t)$ and its complement $C \setminus K_C(t)$ are rectangle-continuous, as explained in example (a) of Subsection 3.4. Moreover, both are rectangles so their geometric loss relative to the family of rectangles is 1. Hence $Loss(K_C, Rectangles) = 1 + 1 = 2$.

(b) Let $C = [0, 1] \times [0, 1]$. Define the following knife function: $K_C(t) = [0, t] \times [0, t] \cup [1 - t, 1] \times [1 - t, 1]$. For every t , $K_C(t)$ is a union of two squares and $C \setminus K_C(t)$ is also clearly seen to be a union of two squares. Since the two squares meet only at their corners, no positive-area squares are created “all of a sudden”; both functions are square-continuous. Moreover, their geometric loss relative to the family of squares is 2. Hence, $Loss(K_C, Squares) = 4$.



$$\begin{aligned} \text{Loss}(K, \text{Rectangles}) &= 2 \\ \text{Loss}(K, \text{Squares}) &= \infty \end{aligned}$$



$$\begin{aligned} \text{Loss}(K, \text{Rectangles}) &= 4 \\ \text{Loss}(K, \text{Squares}) &= 4 \end{aligned}$$




Figure 3: Geometric loss of knife functions. $K(t)$ is filled with horizontal lines. Dotted lines mark future knife locations.

(c) Let $C = R^2$ (the unbounded 2-dimensional plane) and S the family of squares (which, as mentioned in the introduction, also contains the half-planes). Define the following knife function: $K_C(t) = [-\infty, \tan(\pi \cdot (t - \frac{1}{2}))] \times [-\infty, \infty]$. both $K_C(t)$ and its complement $C \setminus K_C(t)$ are half-planes, so their geometric loss relative to S is 1 and the geometric loss of K_C is 2.

4 Envy-Free Division For Two Agents

We now present a generic envy-free division procedure for $n = 2$ agents, based on a knife function.

Claim 2. Let C be a cake, S a family of shapes and $M \geq 2$ an integer. If there is a knife function K_C having $\text{Loss}(K_C, S) \leq M$ Then:

$$\text{Prop}(C, 2, S) \geq \text{Prop}_E(C, 2, S) \geq \frac{1}{M}$$

Proof. C can be divided using the following procedure. (1) Each agent i selects a time $t_i \in [0, 1]$ in which to “stop the knife”. (2) The smaller time is selected; assume it is t_i . (3) Agent i receives $K_C(t_i)$ and agent $1 - i$ receives $C \setminus K_C(t_i)$.

To prove the lower bound, we show that every agent with value measure V can select t_i such that his allocated share P satisfies: (a) $V^S(P) \geq V^S(C \setminus P)$, and (b) $V^S(P) \geq \frac{1}{M}$.

By definition of a knife function, when $t = 0$: $V^S(K_C(t)) = V^S(\emptyset) = 0 \leq V^S(C) = V^S(C \setminus K_C(t))$ and when $t = 1$: $V^S(K_C(t)) = V^S(C) \geq 0 = V^S(\emptyset) = V^S(C \setminus K_C(t))$. The functions $K_C(t)$ and $C \setminus K_C(t)$ are both S-continuous. Hence, by the S-continuity lemma (Subsection 3.4) the functions $V^S(K_C(t))$ and $V^S(C \setminus K_C(t))$ are both continuous functions of t . By the intermediate value theorem, there exists a time t in which the utilities on both sides of the knife are equal: $V^S(K_C(t)) = V^S(C \setminus K_C(t))$. Denote this equal utility by U .

An agent stopping the knife at time t is guaranteed to receive either $K_C(t)$ or a piece that contains $C \setminus K_C(t)$. In both cases the agent feels no envy and has a utility of at least U . Because the geometric loss of the knife function K_C is at most M , by the Chooser Lemma $U \geq \frac{1}{M}$. \square

Corollary 1. By examples (a)-(c) in Subsection 3.6:

(a) $\text{Prop}_E(\text{Rectangle}, 2, \text{Rectangles}) = \frac{1}{2}$, as is already known from classic cake-cutting.

(b) $\text{Prop}_E(\text{Square}, 2, \text{Squares}) \geq \frac{1}{4}$. Combining this with the upper bound of (Segal-Halevi, Hassidim, and Au-

mann 2014), $\text{Prop}(\text{Square}, 2, \text{Squares}) \leq \frac{1}{4}$, gives an equality.

(c) $\text{Prop}_E(\text{Plane}, 2, \text{Squares}) = \frac{1}{2}$.

5 Envy-Free Division For Three Agents

Our procedure for 3 agents is a generalization of the Three Knives procedure (Stromquist 1980). This procedure involves a “sword” moved by a referee, and three “knives” held by the 3 agents. This is an infinite procedure - for every infinitesimal move of the sword, the agents should adjust the locations of their knives. Because the allocated pieces must be connected, no finite algorithm exists (Stromquist 2008), so an infinite procedure is the best that can be hoped for.

To generalize Stromquist’s procedure, we will need two functions: a knife function $K(T)$ on the entire cake, which we call “sword function” (following Stromquist’s terminology); and a family of knife functions $k_T(t)$ on the complement of $K(T)$. We denote the complement of $K(T)$ by $\overline{K(T)} := C \setminus K(T)$ and the complement of $k_T(t)$ by $\overline{k_T(t)} := \overline{K(T)} \setminus k_T(t)$. The following definition extends the geometric loss concept from a single knife function to a pair of a sword function and a family of knife functions:

Definition 7.

$$\text{Loss}(K, k, S) = \max_{T, t \in [0, 1]}$$

$$[\text{Loss}(K(T), S) + \text{Loss}(k_T(t), S) + \text{Loss}(\overline{k_T(t)}, S)]$$

Claim 3. Let C be a cake and S a family of shapes. If there is a sword function $K(T)$ on C and a family of knife functions $k_T(t)$ on $\overline{K(T)}$ and $\text{Loss}(K, k, S) = M$, then: $\text{Prop}(C, 3, S) \geq \text{Prop}_E(C, 3, S) \geq \frac{1}{M}$.

Proof. C can be divided using the following procedure. For every $T \in [0, 1]$:

(1) Each agent i selects a certain $t_i(T) \in [0, 1]$. The selection should be such that every t_i is a continuous function of T (a similar requirement applies in Stromquist’s procedure).

(2) Each agent has the right to shout “cut!”.

(3) If one or more agents shouted “cut!”, then select one of them arbitrarily, call him “the shouter” and give him $K(T)$.

(4) Call the other two agents “the waiters”. Divide $\overline{K(T)}$ between them in the following way:

(4a) Let $t^* = \text{median}_{i=1}^3(t_i(T))$, i.e. the second of the three t_i ’s.




Figure 4: Sword and knife functions of Cor. 2(c). Horizontal lines represent $K(T)$. Vertical lines represent $k_T(t)$.

(4b) Give $k_T(t^*)$ to the waiter with the smaller t_i and give $k_T(t^*)$ to the remaining waiter.

To prove the lower bound, we show that every agent i can select $t_i(T)$ and decide whether to shout “cut”, such that: (a) the utility of his allocated piece is at least as large as the other two pieces, and (b) that utility is at least $\frac{1}{M}$.

In step #1, agent i should select $t_i(T)$ such that $V_i(k_T(t_i)) = V_i(k_T(t_i))$; this is always possible because of the continuity of the knife function k_T . This guarantees that, if agent i is one of the waiters, he will not envy the other waiter.

In step #2, agent i should shout “cut” if he notices that the value of $K(T)$ is equal to the value he would receive in step #4 by not shouting “cut”. Because of the continuity requirements, this guarantees that the shouter will not envy the waiters and vice versa.

By the Chooser Lemma, each agent’s utility is at least $\frac{1}{M}$. \square

Corollary 2. (a) $Prop(Rectangle, 3, rectangles) = \frac{1}{3}$ (a result from classic cake-cutting).

(b) $Prop(Plane, 3, Squares) = \frac{1}{3}$.

(c) $Prop(Square, 3, Squares) \geq \frac{1}{10}$.

Proof. (a) Both the sword function $K(T)$ and the knife function $k_T(t)$ are rectangles growing from the west towards the east. The total geometric loss is clearly $1 + 1 + 1 = 3$.

(b) The sword function is a half-plane bounded by a vertical line moving from $x = -\infty$ to $x = \infty$. The knife function $k_T(t)$ is a quarter-plane to the east of the sword line, bounded by a horizontal line moving from $y = -\infty$ to $y = \infty$. Recall that we treat half-planes and quarter-planes as squares with infinite side-length. Hence the total geometric loss is 3.

(c) The sword function is the union of two corner-squares growing towards the center, as in example (b) of Subsection 3.6. The knife function is the union of four corner-squares as illustrated in Figure 4. Since the squares meet only at their corners, the function is square-continuous. By counting the number of covering squares in different combinations of T and t , it is possible to show that the geometric loss of $k_T(t)$ is at most 4. This is also obviously true for $k_T(t)$. For $K(T)$ the geometric loss is obviously 2. Hence the total geometric loss is 10. \square

6 Conclusion and Future Work

We presented the problem of fairly dividing a cake to two or three agents whose utility functions depend on geometric shape. Our main constructive contributions are generic, symmetric, anonymous division protocols which achieve an envy-free division and a minimum guaranteed utility for every agent.

The tools developed in this paper are generic and can work for cakes and pieces of other geometric shapes. In fact, our tools reduce the envy-free division problem to a geometric problem - the problem of finding appropriate knife functions.

We currently work on extending the results to n agents. This is a challenging task as envy-free division is a difficult problem even for 1-dimensional cakes.

One way to generalize our model is to consider a utility function which takes into account both the value contained in the best square and the total value of the plot, e.g.: $U(P) = WV^S(P) + (1 - W)V(P)$, where W is an agent-dependent constant.

Additional future research topics include: subjective geometric preferences (letting each agent i specify a different family S_i of usable shapes), efficiency and social-welfare maximization.

7 Acknowledgements

Erel Segal-Halevi is grateful to the Doctoral Fellowships of Excellence Program and to the ISF grant 1083/13 for their financial support and Prof. Shmuel Nitzan from the Bar Ilan Economics department for his substantial help in a former version of this paper.

Avinatan Hassidim is supported by ISF grant 1241/12 and by BSF grant 2012344.

Many members of the math.stackexchange.com and mathoverflow.com communities kindly helped in coping with specific geometry-related issues arising in our research. In particular, we are grateful to: Christian Blatter, Ilya Bogdanov, Boris Bukh, Anthony Carapetis, Christopher Culter, David Eppstein, Christopher Fish, Peter Franek, Nick Gill, John Gowers, Dafin Guzman, Marcus Hum, Robert Israel, Barry Johnson, V. Kurchatkin, Raymond Manzoni, Ross Millikan, Mariusz Nowak, Boris Novikov, Joseph O’Rourke, Emanuele Paolini, Steven Taschuk, Marc van Leeuwen, Hagen von Eitzen, Martin von Gagern, Jared Warner, Ittay Weiss and Phoemue X.

References

- Bar-Yehuda, R., and Ben-Hanoch, E. 1996. A linear-time algorithm for covering simple polygons with similar rectangles. *International Journal of Computational Geometry & Applications* 6.01:79–102.
- Barbanel, J. B., and Brams, S. J. 2004. Cake division with minimal cuts: envy-free procedures for three persons, four persons, and beyond. *Mathematical Social Sciences* 48(3):251–269.
- Barbanel, J. B.; Brams, S. J.; and Stromquist, W. 2009. Cutting a pie is not a piece of cake. *American Mathematical Monthly* 116(6):496–514.
- Beck, A. 1987. Constructing a Fair Border. *The American Mathematical Monthly* 94(2):157+.
- Brams, S. J., and Taylor, A. D. 1995. An Envy-Free Cake Division Protocol. *The American Mathematical Monthly* 102(1):9+.
- Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Brams, S. J.; Jones, M.; and Klamler, C. 2008. Proportional pie-cutting. *International Journal of Game Theory* 36(3):353–367.
- Brams, S. J.; Taylor, A. D.; and Zwicker, W. S. 1997. A Moving-Knife Solution to the Four-Person Envy-Free Cake Division. In *Problem. Proceedings of the American Mathematical Society* 125, 547–554.
- Caragiannis, I.; Lai, J. K.; and Procaccia, A. D. 2011. Towards more expressive cake cutting. In *IJCAI*, 127–132. AAAI Press.
- Chen, Y.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2013. Truth, justice, and cake cutting. *Games and Economic Behavior* 77(1):284–297.
- Cohler, Y. J.; Lai, J. K.; Parkes, D. C.; and Procaccia, A. D. 2011. Optimal Envy-Free Cake Cutting. In *AAAI (25th AAAI Conference on Artificial Intelligence)*.
- Dall’Aglia, M., and Maccheroni, F. 2009. Disputed lands. *Games and Economic Behavior* 66(1):57–77.
- Deng, X.; Qi, Q.; and Saberi, A. 2012. Algorithmic Solutions for Envy-Free Cake Cutting. *Operations Research* 60(6):1461–1476.
- Dubins, L. E., and Spanier, E. H. 1961. How to Cut A Cake Fairly. *The American Mathematical Monthly* 68(1):1+.
- Hill, T. P., and Morrison, K. E. 2010. Cutting Cakes Carefully. *The College Mathematics Journal* 281–288.
- Hill, T. P. 1983. Determining a Fair Border. *The American Mathematical Monthly* 90(7):438+.
- Hüsseinov, F., and Sagara, N. 2013. Existence of efficient envy-free allocations of a heterogeneous divisible commodity with nonadditive utilities. *Social Choice and Welfare* 1–18.
- Iyer, K., and Huhns, M. N. 2009. A Procedure for the Allocation of Two-Dimensional Resources in a Multiagent System. *International Journal of Cooperative Information Systems* 18.
- Kurokawa, D.; Lai, J. K.; and Procaccia, A. D. 2013. How to Cut a Cake Before the Party Ends. In *AAAI*.
- Manabe, Y., and Okamoto, T. 2010. Meta-Envy-Free Cake-Cutting Protocols. In Hliněný, P., and Kučera, A., eds., *Mathematical Foundations of Computer Science 2010*, volume 6281 of *Lecture Notes in Computer Science*. Berlin, Heidelberg: Springer Berlin Heidelberg. chapter 44, 501–512.
- Manabe, Y., and Okamoto, T. 2012. A Cryptographic Moving-Knife Cake-Cutting Protocol. *Electronic Proceedings in Theoretical Computer Science* 78:15–23.
- Mirchandani, R. S. 2013. Superadditivity and Subadditivity in Fair Division. *Journal of Mathematics Research* 5(3).
- Reijnierse, J. H., and Potters, J. A. M. 1998. On finding an envy-free Pareto-optimal division. *Mathematical Programming* 83(1-3):291–311.
- Sagara, N., and Vlach, M. 2005. Equity and Efficiency in a Measure Space with Nonadditive Preferences: The Problem of Cake Division. *Journal of Economic Literature* 90.
- Segal-Halevi, E.; Hassidim, A.; and Aumann, Y. 2014. Fair and square: Cake-cutting in two dimensions. *arXiv*.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica* 16(1).
- Stromquist, W. 1980. How to Cut a Cake Fairly. *The American Mathematical Monthly* 87(8):640+.
- Stromquist, W. 2008. Envy-free cake divisions cannot be found by finite protocols. *Electronic Journal of Combinatorics*.
- Su, F. E. 1999. Rental Harmony: Sperner’s Lemma in Fair Division. *The American Mathematical Monthly* 106(10):930+.
- Thomson, W. 2007. Children Crying at Birthday Parties. Why? *Economic Theory* 31(3):501–521.