FAIR DIVISION: CAKE CUTTING, SPERNER'S LEMMA, AND RENTAL HARMONY

PRAKHAR SINGHVI

Abstract. Fair division asks how a resource can be allocated among agents with different preferences. This expository paper studies divisible fair division through the problem of envy-free cake cutting. After developing the general language of valuations, allocations, proportionality, and envy-freeness, the paper examines the classical divide and choose and last diminisher procedures, followed by the Selfridge-Conway procedure for three agents. It then introduces Sperner's lemma and uses it to explain the existence of envy-free cake divisions under natural assumptions on preferences. In the end, the same combinatorial framework is applied to Su's rental harmony theorem, which concerns the assignment of rooms and rents. The paper highlights the contrast between fairness as a simple definition and fair division as a problem involving algorithms, geometry, and combinatorics.

1. Introduction: Fairness, Equality, and Envy

Fair division is the study of how to allocate a resource among several agents who may value that resource differently. The word *fair* is familiar in everyday language, but it does not give it one unique mathematical meaning. In particular, it is important to note that fairness does not simply mean that every agent receives the same physical amount of a resource.

Consider two agents dividing a cake such that one agent values only the chocolate part of the cake, while the other agent values only the vanilla part. A division into equal sized pieces may be very unfair if one agent receives nearly all of the part they value and the other receives almost none. Conversely, an unequal looking division may be fair if each agent receives a piece that they value highly. Thus, the relevant question is not only how large a piece is, but how much the receiving agent values that piece.

This observation leads to one of the central ideas in fair division: "envy-freeness". A division is called *envy-free* if no agent prefers another agent's allocated piece to their own piece. At first glance, envy-freeness seems to be an extremely natural, and perhaps quite simple, requirement. If every single agent prefers their own share to every other share, then no agent has a reason to object to the allocation. The real difficulty is that different agents can value the same piece in very different ways. It is possible that a piece that one agent regards as almost worthless may be the most valuable part of the resource for some other agent.

The classical model for a divisible resource is called *cake cutting*. Here, the *"cake"* can represent land, time, bandwidth, water, or any heterogeneous resource that can be divided into arbitrarily small pieces. Cake cutting provides a clean setting because it allows fairness

Date: July 13, 2026.

conditions to be stated precisely because it supports both explicit algorithms and existence proofs.

The main goal of this paper is to study envy-free division in the divisible setting. We begin with elementary procedures such as divide and choose, which gives an elegant two-agent solution, and the last diminisher method, which guarantees proportionality for any number of agents. We then examine the Selfridge-Conway procedure, which gives an envy-free allocation for three agents. The second half of the paper develops a different approach. Instead of directly constructing a fair division through a sequence of cuts and choices, we use Sperner's lemma to prove the existence of envy-free cake divisions. This approach, developed for fair division by Simmons and explained by Su [5], also leads to the rental harmony theorem for assigning rooms and rents.

The paper is intentionally focused. It follows one central question: How can a divisible resource be allocated so that no agent envies another agent's share? The answer to this question begins with simple cutting procedures, but it eventually leads to combinatorics, geometry, topology, and algorithmic complexity.

The remainder of the paper is organized as follows. Section 2 introduces the general language of agents, resources, valuations, and allocations. Section 3 specializes this language to the cake-cutting model. Section 4 presents divide and choose and the last diminisher method. Section 5 studies the Selfridge-Conway procedure for three agents. Section 6 introduces Sperner's lemma and labelled triangulations. Section 7 applies Sperner's lemma to cake cutting. Section 8 explains rental harmony as a second application of the same combinatorial method. Finally, Section 9 discusses algorithmic issues, and Section 10 gives several directions for further study.

2. The General Model of Fair Division

2.1. **Agents, resources, and feasible shares.**

A fair-division problem begins with a finite set of agents, $N = 1, 2, \ldots, n$, where $n \geq 2$. The agents may be individuals, families, firms, or any entities among whom a resource is to be allocated.

Let X denote the resource being divided, such as a cake, plot of land, collection of rooms, or set of tasks. A *share* is a feasible portion of X assigned to an agent.

The nature of X determines the mathematical model. Some resources, such as books or rooms, are indivisible, while others, such as cake, land, time, or water, are divisible and can be partitioned into small pieces. This paper focuses on divisible resources, though the same basic concepts apply throughout fair division.

2.2. **Valuation functions.**

The central mathematical object in fair division is a valuation function. For each agent $i \in N$, let v_i be the valuation function of agent i. If S is a feasible share of the resource, then $v_i(S)$ denotes the value that agent i assigns to S.

The use of a separate valuation function for every agent is quite essential. Fair division is not merely concerned with the quantity of resource that each agent receives. It is also concerned with how every agent evaluates the share they receive. In general, two agents need not agree about the value of the same share. It is entirely possible that $v_i(S) \neq v_j(S)$ for two distinct agents $i, j \in N$ and the same share S. For example, one agent may value the fruit portion of a cake highly and the chocolate portion very little, while another agent may have exactly the opposite preferences.

For most of this paper, valuations will be nonnegative. Thus, for every agent $i \in N$ and every feasible share S, $v_i(S) \geq 0$. This assumption reflects the fact that cake is a good as receiving more of a valuable piece is never worse than receiving less of it. The theories of chores and mixed goods allow negative values as well, but these settings are outside the main scope of this paper.

2.3. **Allocations.**

An allocation specifies which share each agent receives.

Definition 2.1.
An *allocation* of X among the agents in N is an ordered tuple $A = (A_1, A_2, \ldots, A_n)$, where A_i is the share allocated to agent i.

In standard fair division problem, no part of the resource should be assigned to more than one agent. Therefore, the shares must be pairwise disjoint: $A_i \cap A_j = \varnothing$ for all distinct agents $i, j \in N$. An allocation is called *complete* if every part of the resource is assigned to some agent. In that case, $A_1 \cup A_2 \cup \cdots \cup A_n = X$.

A complete allocation is quite natural model for cake cutting, since one usually wants the entire cake to be distributed. In some variants of fair division, part of the resource may be left unallocated. This possibility is called *free disposal*. We will not use free disposal in the main procedures considered here.

It is important to distinguish an allocation from a valuation. The object A_i is the piece received by agent i, whereas $v_i(A_i)$ is the value that agent i assigns to that piece. A fair allocation must take both objects into account.

2.4. **Proportionality.**

One of the oldest, and most simple, fairness notions is proportionality. Intuitively, if there are n agents, then every agent should receive at least $\frac{1}{n}$th of the total value of the resource according to their own valuation.

We will usually normalize valuations by assuming that $v_i(X) = 1$ for every agent $i \in N$. This does not mean that the agents agree that the resource has one unit of objective value; instead, it helps place each agent's total value on the same scale.

Definition 2.2. A complete allocation $A = (A_1, \ldots, A_n)$ is *proportional* if $v_i(A_i) \geq \frac{1}{n}$ for every agent $i \in N$.

Proportionality is a natural baseline guarantee. It says that no agent believes they have received less than their fair fraction of the whole resource. However, it is important to note that proportionality does not prevent envy. An agent may receive a share worth more than $1/n$ and still prefer another agent's share more than their own share.

2.5. Envy-freeness.

Definition 2.3. An allocation $A = (A_1, \dots, A_n)$ is *envy-free* if $v_i(A_i) \geq v_i(A_j)$ for every ordered pair of agents $i, j \in N$.

In words, agent i does not envy agent j if agent i values their own share at least as highly as agent j's share. An allocation is envy-free iff this condition holds simultaneously for all ordered pairs of agents.

Envy-freeness is stronger than proportionality for complete allocations of divisible goods. The reason is that if agent i does not envy any other agent, then agent i considers their own share at least as valuable as every other share. Since all shares together make up the whole resource, agent i's share must be worth at least the average share according to i's valuation.

Proposition 2.4. *Let $A = (A_1, \dots, A_n)$ be a complete allocation of a resource X. Suppose that each valuation function is additive and normalized by $v_i(X) = 1$. If A is envy-free, then A is proportional.*

Proof.
Fix an agent $i \in N$. Since A is envy-free, we have $v_i(A_i) \geq v_i(A_j)$ for every $j \in N$. Summing these n inequalities gives:

$$nv_i(A_i) \geq \sum_{j=1}^{n} v_i(A_j).$$

Because the allocation is complete and the shares are pairwise disjoint, additivity yields:

$$\sum_{j=1}^{n} v_i(A_j) = v_i(X) = 1.$$

$$\implies {}_i(A_i) \geq 1.$$

Dividing by n, we obtain $v_i(A_i) \geq \frac{1}{n}$.

Since i was arbitrary, the allocation is proportional. ∎

The converse to Proposition 2.4 is false. Consider the following counterexample:

Example 2.5.
Suppose that three pieces P_1, P_2, P_3 are allocated to agents $1, 2, 3$, respectively. Assume that the agents have the following valuations:

	P_1	P_2	P_3
1	$\frac{2}{5}$	$\frac{1}{2}$	$\frac{1}{10}$
2	$\frac{3}{10}$	$\frac{2}{5}$	$\frac{3}{10}$
3	$\frac{3}{10}$	$\frac{3}{10}$	$\frac{2}{5}$

Each row sums to 1, so the valuations are normalized. Under the allocation $A_1 = P_1,\ A_2 = P_2,\ A_3 = P_3$, every agent receives a share worth at least $1/3$ according to their own valuation. Hence the allocation is proportional.

However, $v_1(A_2) = \frac{1}{2} > \frac{2}{5} = v_1(A_1)$. Thus, agent 1 envies agent 2, and the allocation is not envy-free.

Example 2.5 shows why envy-freeness is the more demanding condition. Proportionality guarantees that each agent receives enough value in an absolute sense, whereas envy-freeness compares an agent's allocation directly with every other allocation.

2.6. **Equitability and efficiency.**
Two additional notions help describe the broader landscape of fair division: "Equitability" and "efficiency"

Definition 2.6.
An allocation $A = (A_1, \dots, A_n)$ is *equitable* if $v_1(A_1) = v_2(A_2) = \cdots = v_n(A_n)$.

Equitability asks whether all agents are equally satisfied according to their own valuations. This differs from envy-freeness. Envy-freeness compares one agent's evaluation of different pieces, whereas equitability compares the values that different agents assign to their own pieces.

Definition 2.7.
An allocation $A = (A_1, \dots, A_n)$ is *Pareto optimal* if there is no feasible allocation $B = (B_1, \dots, B_n)$ such that $v_i(B_i) \geq v_i(A_i)$ for every agent $i \in N$, with strict inequality for at least one agent.

Pareto optimality is an efficiency condition rather than a fairness condition. An allocation can be Pareto optimal while still being unfair. Conversely, an allocation can satisfy a fairness condition while failing to use the resource efficiently. One of the recurring themes of fair division is the interaction between these two goals.

3. Divisible Goods and the Cake-Cutting Model

3.1. **The cake as a heterogeneous divisible resource.**
We now specialize the general framework to divisible goods. The standard model represents the cake as the interval $C = [0, 1]$. The point 0 represents the left endpoint of the cake, and the point 1 represents the right endpoint. A subinterval $[a, b]$, where $0 \leq a \leq b \leq 1$, represents a piece of cake.

The use of an interval is just a mathematical simplification. Cutting the cake at a point $x \in [0, 1]$ divides it into the two pieces $[0, x]$ and $[x, 1]$. More generally, if $0 = x_0 \leq x_1 \leq \cdots \leq x_{n-1} \leq x_n = 1$, then the cut points $x_1, \dots, x_{n-1}$ divide the cake into n consecutive pieces: $[x_0, x_1],\ [x_1, x_2],\ \dots,\ [x_{n-1}, x_n]$.

Adjacent closed intervals always share endpoints. In the nonatomic setting considered below, individual points have value zero, so this overlap does not affect valuations. One may instead use half open intervals if a literal partition is required.

The cake is called *heterogeneous* because different parts may have different values to different agents. If every agent valued every point of the cake uniformly, then cutting the cake into equal lengths would be sufficient. The interest of cake cutting comes from the fact that this assumption is generally false.

3.2. **Valuation assumptions in cake cutting.**

For each agent $i \in N$, the valuation function v_i assigns a nonnegative value to each suitable piece of cake. We assume that the valuation functions satisfy several natural properties:

- Valuations are nonnegative: $v_i(S) \geq 0$ for every agent $i \in N$ and every piece $S \subseteq C$.
- Valuations are additive. If S and T are disjoint pieces of cake, then $v_i(S \cup T) = v_i(S) + v_i(T)$. Thus, the value of two nonoverlapping pieces together is the sum of their individual values.
- Valuations are normalized: $v_i(C) = v_i([0, 1]) = 1$ for every agent $i \in N$. This condition allows proportionality to be written as $v_i(A_i) \geq 1/n$.
- The cake is divisible from the perspective of every agent.

Definition 3.1.
A valuation v_i is *nonatomic* if, whenever $0 \leq a \leq b \leq 1$ and $0 \leq \alpha \leq v_i([a, b])$, there exists a point $c \in [a, b]$ such that $v_i([a, c]) = \alpha$.

Nonatomicity is crucial for cake cutting procedures. Divide and choose assumes that an agent can cut the cake into two pieces that they value equally. The last diminisher method assumes that an agent can trim a piece down to a specified value.

A useful special case arises when each agent has a value density. Suppose that agent i has a nonnegative function $f_i : [0, 1] \to \mathbb{R}_{\geq 0}$ such that $\int_0^1 f_i(x)\,dx = 1.$. Then the value of an interval $[a, b]$ is given by $v_i([a, b]) = \int_a^b f_i(x)\,dx$.

The density $f_i(x)$ describes how much agent i values the cake near the point x. Although not every valuation needs to be presented through a density function, this model gives a useful geometric picture of heterogeneous preferences.

3.3. **Connected and disconnected pieces.**

A piece of cake is *connected* if it is a single interval. For example, $\left[\frac{1}{4}, \frac{3}{4}\right]$ is connected. A piece is *disconnected* if it is the union of two or more separated intervals. For example, $\left[0, \frac{1}{4}\right] \cup \left[\frac{3}{4}, 1\right]$ is disconnected.

This distinction matters both mathematically and practically. One practical example is division of land, where receiving one connected plot is usually more useful than receiving several scattered plots. Similarly, when dividing time, a person may prefer one continuous block to several short intervals. However, connectedness makes fair division more difficult. Some envy-free procedures, including the Selfridge-Conway procedure, may allocate disconnected pieces because the trimmings are divided separately from the main pieces.

3.4. **Cake divisions as points of a simplex.**

A division of the cake into n consecutive pieces can be described by the lengths of those pieces. Suppose that the lengths are $x_1, x_2, \ldots, x_n$, where $x_i \geq 0$ for all $i \in N$, and $x_1 + x_2 + \cdots + x_n = 1$.

The vector $(x_1, \ldots, x_n)$ therefore lies in the standard $(n - 1)$-simplex

$\Delta^{n-1} = \{(x_1, \ldots, x_n) \in \mathbb{R}^n : x_i \geq 0 \text{ for all } i \in N,\ \sum_{i=1}^n x_i = 1\}$.

For two agents, Δ^1 is a line segment. A point $x \in [0, 1]$ represents a division in which first piece has length x and the second piece has length $1 - x$. For three agents, Δ^2 is a triangle. Each point inside the triangle represent a division of the cake into three consecutive pieces whose lengths sum to 1.

This geometric viewpoint will become central in Sections 6 and 7. Sperner's lemma applies to triangulations of simplexes, and the simplex Δ^{n-1} provides the natural space of possible cake divisions.

4. Elementary Cake-Cutting Procedures

4.1. **Divide and choose.**

The divide and choose procedure is the classical method for two agents.

1. Agent 1 cuts the cake into two pieces that they value equally.
2. Agent 2 chooses one of the two pieces.
3. Agent 1 receives the remaining piece.

The procedure is strikingly simple, but it already captures the central idea of fair division. The agent who cuts is protected because they regard the two pieces as equal. The agent who chooses is protected because they select the piece they prefer.

Theorem 4.1. *The divide and choose procedure produces an envy-free, and hence proportional, allocation for two agents.*

Proof. Agent 1 cuts the cake into two pieces P and Q such that $v_1(P) = v_1(Q) = \frac{1}{2}$. The allocation is envy-free because agent 2 chooses the piece they prefer, so agent 2 cannot envy agent 1 and agent 1 values P and Q equally, so agent 1 cannot envy agent 2. Proposition 2.4 then implies that the allocation is proportional. □

The procedure is asymmetric in its roles but symmetric in its guarantee. Each agent obtains a piece they value at least as much as the other piece.

4.2. **The last diminisher method.**

Divide and choose does not directly extend to more than two agents. One classical procedure for n agents is the last diminisher method. This method guarantees proportionality, although it does not generally guarantee envy-freeness.

Suppose that n agents are dividing a cake. Agent 1 cuts off a piece that they value as exactly $1/n$. The other agents inspect the proposed piece in turn. Whenever an agent believes that the proposed piece is worth more than $1/n$, that agent trims it down to a piece worth exactly $1/n$ according to their own valuation. Then, the trimmings are returned to the remainder of the cake. After every agent got the opportunity to trim, the last agent who trimmed the piece receives it. If nobody trims, agent 1 simply receives it. The remaining $n - 1$ agents repeat the procedure on the remaining cake.

Theorem 4.2. *For nonatomic additive valuations, the last diminisher method produces a proportional allocation for any number $n \geq 2$ of agents.*

Proof.
We proceed by induction on n. The case $n = 2$ is immediate, since the method reduces to a proportional two-agent division.

Assume that the statement holds for $n-1$ agents, and consider an instance with n agents. Let P be the final proposed piece after all possible trimmings, and let agent k be its recipient. By construction, $v_k(P) = \frac{1}{n}$.

Now fix any remaining agent $i \neq k$. When agent i examined the proposed piece, either they trimmed it to value $1/n$, or they regarded it as worth at most $1/n$. Subsequent trimmings can only decrease the piece. Therefore, $v_i(P) \leq \frac{1}{n}$.

Let $R = C \setminus P$ be the remaining cake. Since $v_i(C) = 1$, we have $v_i(R) = 1 - v_i(P) \geq \frac{n-1}{n}$.

By the induction hypothesis, the remaining $n-1$ agents can divide R proportionally relative to R. Thus, agent i receives a piece $A_i \subseteq R$ satisfying:

$$v_i(A_i) \geq \frac{1}{n-1} v_i(R) \geq \frac{1}{n-1} \cdot \frac{n-1}{n} = \frac{1}{n}.$$

The recipient k also receives value exactly $1/n$. Hence every agent receives at least $1/n$ according to their own valuation. □

The last diminisher method illustrates a basic contrast. Proportionality is often attainable through a recursive argument because every agent only needs a guaranteed fraction of their own total value. Envy-freeness is more hard to achieve as it requires every agent to compare their own final piece with every other final piece.

5. The Selfridge-Conway Procedure for Three Agents

The divide and choose procedure gives an envy-free allocation for two agents. The next natural question that arises here is whether there is an equally explicit envy-free procedure for three agents. The answer is yes, but the procedure is much more intricate. The classical Selfridge-Conway procedure, described in detail by Brams and Taylor [2], gives an envy-free division for three agents.

Let the three agents be called A, B, and C. The procedure has two stages. First, the main cake is divided into three pieces. Next, if a trimming was created during the first stage, that trimming is divided separately.

5.1. **The main pieces.**

1. Agent A cuts the cake into three pieces that they value equally.
2. Agent B examines the three pieces. If B has a unique favorite piece, then B trims that piece so that the trimmed main piece is tied, according to B, with B's second favorite piece. The removed part is called the *trimmings*. If B does not have a unique favorite, then no trimming is necessary.
3. Agent C chooses a main piece first.
4. Agent B chooses next. If C did not choose the trimmed main piece, then B must choose the trimmed main piece.

5. Agent A receives the remaining main piece.

If no trimming occurred, then the procedure ends there. Because agent A views all three pieces as equal, agent B receives one of their favorite remaining pieces, and agent C chooses first; thus, no agent envies another.

Suppose now that a trimming occurred. Let T denote the agent among B and C who received the trimmed main piece, and let D denote the other agent among B and C. The trimmings are divided as follows.

5.2. **The trimmings.**

1. Agent D divides the trimmings into three pieces that they value equally.
2. The agents choose pieces of the trimmings in the order T, then A, then D.

The final allocation may give the holder of the trimmed main piece a disconnected share: the trimmed main piece together with one part of the trimmings. This is one reason why connectedness is a genuine additional constraint in cake cutting.

Theorem 5.1. *For three agents with nonatomic additive valuations, the Selfridge-Conway procedure produces an envy-free allocation.*

Proof.
We verify that each agent does not envy another.

First consider the main pieces. Agent C chooses first; hence, agent C does not envy another agent's main piece. Agent B is protected by the trimming rule. If C chooses the trimmed main piece, then B will choose an untrimmed piece which B values at least as highly as the trimmed piece. If C chooses an untrimmed piece, then B will receive the trimmed main piece, which B made equal in value to their second favorite original piece. In either case, B does not envy another agent's main piece. Agent A cut the original three pieces to have equal value. Also, A does not receive the trimmed main piece; hence, A receives an untrimmed piece that A regards as at least as valuable as the trimmed main piece.

Still, it remains to consider the trimmings. Agent D divides the trimmings into three equal pieces according to their own valuation. Hence D does not envy another agent's trimming piece. Agent T chooses first among the three trimming pieces, so T does not envy another agent's trimming piece. Agent A chooses second and therefore receives a trimming piece at least as valuable, according to A, as the trimming piece received by D.

We now compare the complete bundles. Two agents A and D receive an untrimmed main piece, and A values those two main pieces equally. Since A chooses a trimming piece before D, agent A does not envy D. The only agent who receives a trimmed main piece is T. Agent A values the original untrimmed version of that main piece as highly as their own main piece. The bundle of T consists of that trimmed main piece together with only part of the trimmings. The original untrimmed main piece consists of the trimmed main piece together with all of the trimmings. And hence, A does not envy T.

Ultimately, T does not envy another complete bundle as T does not envy another main piece and chooses first among the trimming pieces. Likewise, D does not envy another complete bundle because D does not envy another main piece and regards all three trimming pieces as equal. Thus, none of the three agents envies another, and the final allocation becomes envy-free. $\square$

The proof demonstrates the significant role of trimming. Agent B identifies the only main piece that could create serious envy and reduces its value to a controlled level. The later division of the trimmings removes the remaining possible source of envy. This procedure is much more complicated than divide and choose, but every step has a precise purpose.

6. Sperner's Lemma: The Combinatorial Engine

The procedures in the previous section construct fair divisions by asking agents to cut, trim, and choose. We now turn to a different method. Sperner's lemma is a combinatorial theorem about labelled triangulations of simplexes. Despite its elementary statement, it has powerful consequences in topology and fair division. Sperner introduced the lemma in 1928 [4]; its use in the present fair-division setting is developed by Su [5].

6.1. Simplexes and triangulations.

A 0-simplex is just a point, a 1-simplex is a line segment, a 2-simplex is a triangle, and a 3-simplex is a tetrahedron. More generally, a d-simplex is the convex hull of $d+1$ affinely independent points.

A *triangulation* of a d-simplex is a subdivision of it into smaller d-simplexes, which are called *elementary simplexes*, such that any two elementary simplexes either do not meet or meet in a common face. The largest distance between two points in one elementary simplex is known as the *mesh* of the triangulation. A fine triangulation has a small mesh.

6.2. **Sperner labellings.** Let S be a d-simplex whose main vertices are labelled $1, 2, \ldots, d+1$. Each facet of S, which means each $(d-1)$-dimensional face, is opposite one of the main vertices. The facet opposite the main vertex labelled j is called facet j.

Definition 6.1.
A labelling of the vertices of a triangulation of S by $1, 2, \ldots, d+1$ is a *Sperner labelling* if no vertex on facet j receives label j.

Equivalently, every vertex lying on a face spanned by some collection of main vertices may use only the labels of those main vertices. An elementary d-simplex is *fully labelled* if its $d+1$ vertices have all distinct labels $1, 2, \ldots, d+1$.
Figure 1 illustrates the two-dimensional case. The vertices of the large triangle are labelled 1, 2, and 3. Along the edge between vertices 1 and 2, only labels 1 and 2 may occur. Similar restrictions hold on the other two edges. Interior labels are unrestricted.

Theorem 6.2 (Sperner's lemma)**.**
Every Sperner-labelled triangulation of a d-simplex contains an odd number of fully labelled elementary d-simplexes. In particular, it contains at least one fully labelled elementary simplex.

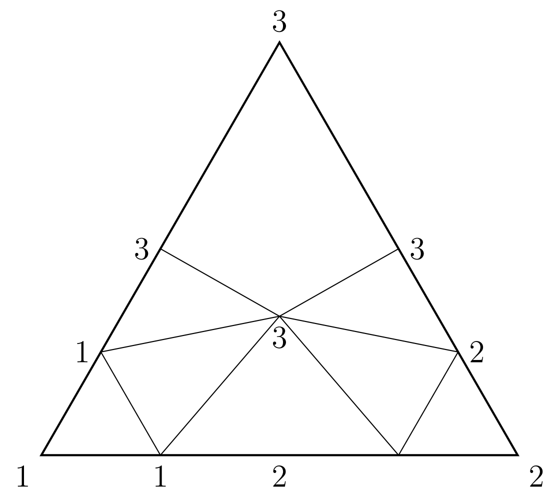

Figure 1. A labelled triangulation of a triangle. On each boundary edge, only the labels of that edge's endpoints appear.

Proof.
We prove the theorem by induction on d. For $d = 1$, the simplex is a line segment with endpoints labelled as 1 and 2. As one moves from the first endpoint to the second, the label must change odd number of times; hence, there are an odd number of elementary edges with labels 1 and 2.

Assume that the theorem holds in dimension $d - 1$. Now, consider a Sperner-labelled triangulation of a d-simplex S. Call a $(d-1)$-face of an elementary simplex a *door* if its vertices carry the labels $1, 2, \ldots, d$. By Sperner condition, only facet $d + 1$ of the large simplex can contain the boundary doors. The induced labelling on this facet is Sperner labelling in dimension $d - 1$. So the induction hypothesis implies that the number of boundary doors is odd.

By inspectng an elementary d-simplex, it becomes clear that if it has a door and is fully labelled, then it has exactly one door, namely the face obtained by omitting the vertex labelled $d + 1$. If it has a door but is not fully labelled, then its labels consist of $1, 2, \ldots, d$ together with one repeated label among these d labels. In that case it has two doors exactly. Hence, modulo 2, the number of doors of an elementary simplex is 1 precisely when the simplex is fully labelled.

Every interior door belongs to exactly two elementary d-simplexes. Hence, contributes an even amount to the total count of doors. Therefore, modulo 2, the number of fully labelled elementary d-simplexes has the same parity as the number of boundary doors. As the latter number is odd, the number of fully labelled elementary d-simplexes is odd as well. □

The proof has quite useful algorithmic interpretation. Starting from a boundary door, we may move into the adjacent elementary simplex. If that simplex is fully labelled, the search ends. Otherwise, the simplex has exactly one other door, which determines the next step. This is refered to as the *trap-door path*. As the triangulation is finite, the path eventually reaches a fully labelled simplex or another boundary door. Then, the parity argument guarantees that at least one boundary path reaches a fully labelled simplex.

7. Sperner's Lemma in Cake Cutting

Sperner's lemma gives a conceptually different proof that envy-free cake divisions exist. The approach explained in this section follows the framework presented by Su [5]. Rather than having agents just perform a fixed sequence of cuts and choices, we study the entire space of possible cut-sets at once.

7.1. **Preferences over cut-sets.**

Fix n agents. A point $x = (x_1, \dots, x_n) \in \Delta^{n-1}$ represents a cut-set that divides the cake into n consecutive pieces of physical lengths $x_1, \dots, x_n$. The pieces themselves depend on the cumulative cut positions:

$$x_1, \quad x_1 + x_2, \quad \dots, \quad x_1 + \cdots + x_{n-1}.$$

For an agent $i \in N$, let $P_i(x) \subseteq N$ denote the set of pieces that agent i weakly prefers under the cut-set x. Thus, $j \in P_i(x)$ means that agent i regards piece j as at least as valuable as every other piece in that division.

We now make two assumptions about preferences:

(1) *Hunger.* If a piece has zero physical length, then no agent prefers it to a piece with positive length.
(2) *Closed preference sets.* For every agent $i \in N$ and every piece $j \in N$, the set $\{x \in \Delta^{n-1} : j \in P_i(x)\}$ is closed.

The hunger assumption prevents the situation when an empty piece is chosen. The closedness assumption confirms that a preference which persists along a convergent sequence of cut-sets also persists at the limiting cut-set. For continuous value densities, these assumptions are quite natural.

Theorem 7.1. *Suppose that n agents have hungry, closed preference sets over the consecutive pieces of a cake. Then there exists a cut-set in which the agents can be assigned distinct preferred pieces. Consequently, an envy-free cake division exists.*

7.2. **Ownership labels.**

To apply Sperner's lemma, triangulate the simplex Δ^{n-1} very finely. So, we need a way to associate each vertex of the triangulation with one of the n agents such that every elementary $(n-1)$-simplex has exactly one vertex owned by each agent.

Barycentric subdivision provides such an ownership rule. In one barycentric subdivision, every elementary simplex is determined by a chain of faces of dimensions $0, 1, \dots, n-1$. Its vertices are the barycenters of those faces. We may assign owner i to the barycenter of a face of dimension $i - 1$. Thus, every elementary simplex has one vertex owned by each agent. Repeating the subdivision makes the mesh arbitrarily small while preserving this ownership structure.

At every vertex z of the triangulation, ask the owner of z which cake piece they prefer under the cut-set, being represented by z. Label z by the number of a preferred piece. If the owner has several preferred pieces, simply choose any one of them.

7.3. **Why the preference labels form a Sperner labelling.**

The resulting labels satisfy the boundary condition of Sperner's lemma. Suppose that a point $z = (x_1, \ldots, x_n)$ lies on the facet of Δ^{n-1} where $x_j = 0$. At this cut-set, piece j is empty. By hunger, the owner of z cannot prefer piece j. Hence z cannot receive label j.

Therefore, the preference labels form a Sperner labelling. By Theorem 6.2, there exists a fully labelled elementary $(n-1)$-simplex. As every elementary simplex also has one vertex owned by each agent, this fully labelled simplex represents n nearby cut-sets in which different agents prefer different pieces.

7.4. **Passing to an exact division.**

We now prove Theorem 7.1.

Proof of Theorem 7.1.
For each positive integer m, choose a triangulation of Δ^{n-1} with mesh less than $1/m$, together with an ownership labelling as above. Sperner's lemma gives a fully labelled elementary simplex σ_m.

For every agent $i \in N$, let $z_{m,i}$ be the unique vertex of σ_m owned by agent i. Because σ_m is fully labelled, the preferred-piece labels on the vertices $z_{m,1}, \ldots, z_{m,n}$ are all distinct. Thus, they define a permutation π_m of N, where agent i prefers piece $\pi_m(i)$ at $z_{m,i}$.

There are only finitely many permutations of N. Hence, after passing to a subsequence, we may assume that $\pi_m = \pi$ is independent of m. Choose an arbitrary point $q_m \in \sigma_m$. Since Δ^{n-1} is compact, a further subsequence of (q_m) converges to some point $q \in \Delta^{n-1}$.

The mesh of σ_m tends to zero. Therefore, for every agent $i \in N$, the vertex $z_{m,i}$ also converges to q. Agent i prefers piece $\pi(i)$ at each cut-set $z_{m,i}$. By closedness of the preference set for piece $\pi(i)$, agent i prefers piece $\pi(i)$ at the limiting cut-set q. Since π is a permutation, the agents prefer distinct pieces at q. Assigning agent i the piece $\pi(i)$ gives an envy-free division. □

This argument is independent of the explicit computation of the cut-set. The compactness argument shows that the envy-free cut-set exists. The trap-door access to Sperner's lemma provides a finite procedure for finding a fully labelled simplex in any finite triangulation of the simplex. Hence, this provides an approximate algorithm that becomes more accurate as the meshrefinement increases.

Remark 7.2. The approximation obtained directly from a fine triangulation is actually geometric as the cut-sets at the vertices of a fully labelled elementary simplex are physically close. To translate this into a numerical bound on the agents' possible envy, we need a additional continuity condition on valuations. Under continuous value densities, sufficiently small changes in cut positions produce sufficiently small changes in values.

8. Rental Harmony: Sperner's Lemma Beyond Cake Cutting

Cake cutting concerns a divisible good. Rental harmony combines two different kinds of objects: rooms are indivisible, while rent is divisible. Suppose n housemates must choose among n rooms, while they also must divide a fixed total rent. Now, here the problem is to assign one room to each housemate and set a price for every room so that no tenant prefers another tenant's room at its assigned price.

This problem is called *rental harmony*. The key point is that a room cannot be split, but its price can be adjusted. Su showed that Sperner's lemma still gives an existence proof under mild assumptions [5].

8.1. **The price simplex.**

Normalize the total rent to 1. A pricing scheme is then just a vector $p = (p_1, \dots, p_n) \in \Delta^{n-1}$, where $p_j \geq 0$ is the price of room j and $p_1 + p_2 + \cdots + p_n = 1$.

Thus, the set of all possible price divisions is again the simplex Δ^{n-1}. This is the first structural similarity between cake cutting and rent division.

For a tenant $i \in N$, let $R_i(p) \subseteq N$ be the set of rooms that tenant i weakly prefers at the pricing scheme p. The preferences here may be arbitrary. Specifically, we do not assume that they are given by a formula such as room value minus price.

8.2. **The rental harmony theorem.**

Theorem 8.1 (Rental Harmony Theorem)**.** Suppose that n housemates must divide a total rent among n rooms. Assume the following conditions:

(1) *Good house.* At every pricing scheme, each housemate finds at least one room acceptable.
(2) *Miserly tenants.* Whenever at least one room is free, each housemate prefers a free room to every nonfree room.
(3) *Closed preference sets.* For every tenant and room, the set of pricing schemes at which that room is preferred is closed.

Then there exists a pricing scheme and a assignment of one room to each housemate such that every housemate weakly prefers their assigned room at its assigned price to every other room at its assigned price.

The conclusion says exactly that the room and rent allocation is envy-free.

8.3. **Why cake cutting and rent division differ.**

At first sight, it may appear that the cake-cutting proof applies immediately. At every vertex of a triangulation of the price simplex, assign an owner and ask that owner which room they prefer. The good house condition ensures that some room can be chosen.

However, the resulting labelling is not a standard Sperner labelling. In cake cutting, a point on the facet $x_j = 0$ represents a division in which cake piece j is empty, and so hunger prevents label j. In rent division, a point on the facet $p_j = 0$ represents a pricing scheme in which room j is free. The miserly tenants condition makes the label j especially likely rather than impossible.

This reversal is central obstacle. The condition on a boundary face is opposite of usual Sperner condition.

8.4. **Dualization.**

Su resolves this obstacle by using a dual complex of the triangulated simplex. Dualization reverses the roles of faces of complementary dimensions. A vertex of the original simplex corresponds to a facet in the dual complex, while a facet of the original simplex corresponds to a vertex in the dual complex.

The original rent-preference labelling has the following property: on a face where several rooms are free, every owner chooses one of the free rooms. As soon as passed to the dual complex, this becomes precisely the boundary restriction required for a Sperner labelling. Sperner's lemma therefore produces a fully labelled elementary simplex in the dual complex, which corresponds to a fully labelled elementary simplex in the original triangulation.

A fully labelled simplex means that nearby pricing schemes assign different preferred rooms to different tenants. By using a sequence of finer and finer triangulations, compactness and closed preference sets yield a limiting price vector at which the tenants prefer distinct rooms. This proves Theorem 8.1; see Su [5, Sections 6–7] for the full dualization argument.

The rental harmony theorem is valuable not only because it solves a familiar roommate problem. It shows that the simplex-and-labelling framework is not tied to cake cutting. In cake cutting, coordinates of the simplex describe lengths of cake pieces. In rental harmony, coordinates describe prices of rooms. In both cases, Sperner's lemma turns local preference information into a global fair-division conclusion.

9. Algorithmic Difficulty and Further Directions

The preceding results raise an important question. It is one thing to prove that an envy-free division exists. It is another to find such a division efficiently using a reasonable number of questions about agents' preferences.

9.1. **The Robertson-Webb query model.**

A standard formal model for cake-cutting algorithms is the Robertson-Webb query model [3]. An algorithm does not receive the entire valuation function of each agent at once. Instead, it may ask two types of questions.

An *evaluation query* asks agent i to report the value of an interval: $\mathrm{Eval}_i(a,b) = v_i([a,b])$.

A *cut query* asks agent i to identify a point b such that a piece beginning at a has a specified value α: $\mathrm{Cut}_i(a,\alpha) = b \quad \text{where} \quad v_i([a,b]) = \alpha$.

Divide and choose and last diminisher can be interpreted in this model. The query model makes it possible to measure the computational cost of a cake-cutting procedure by counting how many preference questions it asks.

9.2. **Existence, boundedness, and practicality.**

The Selfridge-Conway procedure is a short bounded protocol for three agents. The general case is far more difficult. Aziz and Mackenzie gave a discrete bounded envy-free protocol for any number of agents [1]. Their result proves that there is a finite bound, depending only on n, on the number of queries required to find an envy-free allocation. The known bound is extraordinarily large, which emphasizes the difference between theoretical boundedness and practical usability.

The Sperner approach yields a different kind of constructive insight. Given a finite triangulation, the trap-door search finds a fully labelled simplex by following a combinatorial path. As the triangulation becomes finer, the resulting divisions approximate an exact envy-free division ever more closely. The method is conceptually elegant, but the number of vertices in a fine high-dimensional triangulation grows rapidly.
These observations illustrate a recurring distinction in fair division:
existence $\neq$ constructibility $\neq$ efficient computation.

An allocation may exist by a compactness argument, be approximable by a finite search, and still be difficult to compute in practice.

9.3. **Connected pieces and extensions.**

Several natural extensions of the cake-cutting problem remain challenging. One may require every agent to receive a connected interval rather than a disconnected union of intervals. One may also ask for fair division of two-dimensional land, for divisions with unequal entitlements, or for procedures that remain fair when agents may strategically misrepresent their preferences.

Another direction is to replace goods by chores. Here lower cost is preferable to higher cost, so the inequalities defining envy reverse their interpretation. Su's framework also applies to chore division by viewing the allocation of rent payments as an allocation of undesirable burdens [5]. Such extensions show that the central ideas of fair division are robust, even though the exact mathematical conditions change from one setting to another.

10. Further Questions

The methods developed in this paper suggest several questions that lead beyond the present exposition.

First, can envy-free cake-cutting procedures be made simpler and more efficient for many agents? The existence of bounded protocols is a major theoretical achievement, but a procedure with an enormous worst-case query bound is not a practical roommate tool. Finding transparent procedures with substantially lower complexity remains an important goal.

Second, what additional difficulty is caused by connectedness? The Selfridge-Conway procedure may give disconnected pieces because the trimmings are handled separately. In many applications, however, an agent must receive one contiguous piece of land, one continuous interval of time, or one physically usable region. Fairness and geometric constraints must then be considered together.

Third, the rental harmony theorem invites a broader question. Which preference assumptions are genuinely necessary to guarantee an envy-free division of rooms and rents? The miserly-tenants condition is natural when tenants strongly prefer free rooms, but real preferences may be more complicated. Understanding the weakest useful hypotheses is part of the mathematical content of the problem, not merely a technical detail.

The study of fair division begins with a familiar human concern: how can people divide something without feeling cheated? The answer is not simply to divide the resource into equal physical parts. Instead, the mathematics must account for subjective values, compare several competing notions of fairness, and construct allocations that respect those values. Cake cutting, Sperner's lemma, and rental harmony show how an ordinary question about sharing can lead to rich interactions among algorithms, combinatorics, and geometry.

References

[1] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science*, pages 416–427. IEEE, 2016.
[2] Steven J. Brams and Alan D. Taylor. An envy-free cake division protocol. *American Mathematical Monthly*, 102(1):9–18, 1995.
[3] Jack M. Robertson and William A. Webb. *Cake-Cutting Algorithms: Be Fair If You Can.* A K Peters, Natick, Massachusetts, 1998.
[4] Emanuel Sperner. Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 6:265–272, 1928.
[5] Francis Edward Su. Rental harmony: Sperner's lemma in fair division. *American Mathematical Monthly*, 106(10):930–942, 1999.