Finding Envy-Free Divisions

Ari Seeger

July 2024

Abstract

The problem of envy-free division, where resources are divided among individuals such that no one envies another's allocation, has significant applications in fair division and resource allocation. This paper explores various mathematical and topological approaches to finding, or disproving, the existence of an envy-free division. We investigate convex combinations of divisions to derive new envy-free solutions and analyze the construction of polytopes and their coverings to demonstrate the existence or non-existence of envy-free divisions. Additionally, we utilize Sperner's Lemma and the KKM Lemma to provide an alternative method for identifying envy-free divisions. These solutions for finding fair divisions can be applied to everyday situations, allowing players to avoid envy and discontent with their allocations.

1 Introduction

The concept of fair division is a fundamental problem in the fields of mathematics, economics, and game theory. It involves dividing a set of resources among multiple individuals in the best possible way so that everyone receives an amount they believe is deserved. This principle can be used for ensuring fairness and equity in various real-world scenarios, such as inheritance division, resource allocation, and collaborative tasks.

Ensuring an envy-free division is not only theoretically interesting but also practically valuable. Envy-free divisions are fair divisions in which each player believes they got the highest valued allocation. Envy-free fair division has applications in game theory and almost any system involving a splitting of resources. Despite its importance, finding envy-free solutions can be challenging, especially as the number of participants and the complexity of the resources increase.

This paper aims to explore multiple mathematical and geometric approaches to achieving envy-free divisions. We will examine the use of convex combinations of divisions to discover new envy-free solutions and construct polytopes with coverings to visualise divisions and the existence of envy-free partitions. Furthermore, we will use Sperner and KKM Lemmas as methods for identifying envy-free divisions in cake-cutting problems involving several players.

This paper begins with algorithms and existing methods for envy-free division. We then discuss approaches using convex combinations of divisions. Section 4 delves into the construction of polytopes and their coverings. Section 5 presents the application of Sperner's Lemma to the problem. Finally, in Section 6, we further our understanding of poltopes and coverings with the KKM Lemma and the Colorful KKM Theorem to map out possible envy-free divisions.

2 Simple Envy-Free Algorithms

Solutions to cake-cutting fair division problems can range from basic two-player methods to complex algorithms used to find fair divisions of multiple cakes and several players. When a single homogenous cake is being divided among n players, a simple and obvious solution is dividing the cake into n slices. However when multiple different cakes are involved and players have varying preferences, divisions can be found which leaves players with a piece they value as greater than $\frac{1}{n}$.

One great example of this is the divider-chooser method for two players. The first player divides the cake into two pieces that they value as equal, and the second player chooses the slice they believe is greater. This division is envy-free, as the first player believes both slices are identical, while the second player believes they got a greater slice.

Though this method works for two players, it is not efficient and leaves each player with very different valuations of their received portion. Another issue with this method is that it does not generalize well for larger numbers of players. However, a variation exists that does accomodate for three people.

The first player can divide the cake into three equally sized pieces. We can label these pieces as A, B, and C. The second player then picks the largest piece and cuts it so it is equal in value to the second largest piece. Suppose he chooses slice A and trims it down with the part trimmed off labeled as $A1$. Now, the third player is free to choose any piece. If the third player does not pick slice A, then the second player must take it. The first player then takes the remaining slice. Now, all that is left to be divided is $A1$, the trimmed parts. We know that slice A was taken either by player two or player three, so the one that did not take it will divide $A1$ into three equal parts. The player with slice A then chooses first, player one chooses second, and the player who divided $A1$ takes the last piece.

We can now examine this division and see if it is envy-free. From player one's perspective, the person with slice A got less than a third of the cake because it was cut smaller. He believes slices B and C to be equal in size, and because he chose his part of $A1$ before the third player, his total is greater. From player two's perspective, he also got the largest piece because he chose one of the two larger pieces to begin with, and then chose first when picking out of $A1$. For player three, he also got the largest because he chose the largest of the pieces A, B, and C, and for the remaining piece $A1$ he divided it equally.

This method creates an envy free division for three players and a single cake, but it can't be generalized. There is however a fair division algorithm which applies to any number of players.

The lone divider is performed with one designated divider who divides the cake into n pieces of equal size. Each player then declares which piece they want most. If multiple players select the same slice, then the divider takes a slice which was not chosen, and every other slice which was only picked by one player is given to the player who chose it. This process is then repeated for the rest of the players with the remaining cake.

Though the lone divider method may seem like it works, it is not envy-free. Each player receives a portion they value as at least $\frac{1}{n}$, but it is possible that they believe another player received more.

To get a better understanding of envy-freeness, we can define all our values with variables.

Given a cake C to be divided among n players, each player i has:

- A value V_i on slices or subsets of C.
- A weight w_i that represents the fraction of C the player is entitled to.

All w_i should add up to 1. We want to divide C into n parts so that for any two players i and h:

$$\frac{V_i(X_i)}{w_i} \geq \frac{V_i(X_h)}{w_h}$$

This ensures that player i does not envy player h, considering their different entitlements and player i's perspective of the value.

We can use this equation to determine if divisions are envy-free by applying it to all players. For example, suppose we have three players and three cakes. Player one values chocolate three times more than vanilla, and strawberry two times more than vanilla. Player two values both vanilla and chocolate as twice as good as strawberry. Player three values them all equally. One possible division leaves player one with two thirds of the chocolate cake and one third of the strawberry while player two gets one third of the chocolate and one two thirds of the vanilla. Player three gets two thirds of the strawberry and one third of the vanilla. Let's now calculate the values from each person's perspective. We can represent amounts as percentages of the whole according to each player.

Player	**Player 1's Portion**	**Player 2's Portion**	**Player 3's Portion**
Player 1	44%	28%	28%
Player 2	33%	40%	27%
Player 3	33%	33%	33%

Values from each player's perspective

As we can see from the table, each payer believes they got a portion equal or better than the portions of everyone else. We can use this equation to continue calculating values from the perspectives of players.

3 Convex Combinations

Another solution to fair division problems involves using money, or some other currency to compensate for smaller portions in a division. A player who receives a larger slice of the cake may have to pay a certain amount while someone with less of the cake would receive money to make up for it. This method can create more solutions to a problem which may have only had one.

Consider a problem with 3 players dividing some object. For this example the object is homogenous and can be divided or cut into parts of any size, such as a cake or a large collection of smaller objects. Each player values the cake differently, and they all feel entitled to at least $\frac{1}{3}$ of the entire cake. Say Player A values the cake at \$30, Player B believes the cake is worth \$24, and Player C thinks the value of the cake is only \$18. One possible division is where Players A and C each take half the cake. Player A gives \$5 away, Player C gives \$3 away, and Player B receives \$8. With this division, each player receives what they believe to be $\frac{1}{3}$ of the cake. Player A values the division he received as $\frac{1}{2}(\$30) - \$5 = \$10$, where $\frac{1}{2}$ is the portion of the cake, \$30 is his valuation of the cake, and \$5 is the amount he gave away. The total amount he received is then \$10 out of the \$30 cake, or $\frac{1}{3}$. The same is true for Player C, who values his division as $\frac{1}{2}(\$18) - \$3 = \$6$. As for Player B, he received \$8, which is $\frac{1}{3}$ of his valuation for the cake which was \$24. However, this division is not envy-free. Even though each player believes they got $\frac{1}{3}$ of the cake, they each also believe another player got more than they did. Player A is envious of Player C because he thinks Player C got $\frac{1}{2}(\$30) - \$3 = \$12$. Player B is envious of Player C as well because he thinks Player C got $\frac{1}{2}(\$24) - \$3 = \$9$, which is more than the \$8 he received. Finally, Player C is envious of Player B because Player B received \$8 which is more than he values his division as. If we are searching for envy-free divisions, then one option is to split the cake into thirds with no money given. This way every player values their cake as equal to everyone else's. Another option is to give the entire cake to Player A, who values it highest, and have him pay each other player what he believes is $\frac{1}{3}$. Because the cake is worth most to him, his value for $\frac{1}{3}$ of the cake is higher than the others, thus they each receive more than $\frac{1}{3}$ from their perspective. These divisions can be represented in a chart:

Size of Portion (A, B, C)	Money Received (A, B, C)
$(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$	$(0, 0, 0)$
$(1, 0, 0)$	$(-20, 10, 10)$

With at least two envy-free divisions, another can be found using convex combinations [5]. To understand how this works, imagine the cake is split into two or more different parts to begin with. With the three players, they can perform a different envy-free division with each of those parts of the cake, and once everyone's money and slices are put together, we have found a new envy-free division for the whole cake.

For example, if $\frac{1}{2}$ of each division is used, then the division of cake is found by multiplying each by $\frac{1}{2}$ and adding them:

$$\frac{1}{2}(\frac{1}{3},\frac{1}{3},\frac{1}{3})+\frac{1}{2}(1,0,0)=(\frac{2}{3},\frac{1}{6},\frac{1}{6})$$

This is repeated for the money vectors:

$$\frac{1}{2}(0,0,0)+\frac{1}{2}(-20,10,10)=(-10,5,5)$$

Calculating the values received for each player shows that this division is also envy-free. From the perspective of Player A, each player received \$10. According to Player B, he and Player C each received \$9 while Player A only got \$6. For Player C, they both got \$7 while Player A lost \$2.

Using convex combinations, there are unlimited envy-free divisions that can be found. However, envy-free divisions do not guarantee that every player receives the most value. The most obvious and simple division would be to split the cake into thirds, giving each slice to a different player. This is an envy-free division, as each player receives exactly the same thing, but it is not pareto-efficient, meaning there is another way the cake can be divided in which each player is either in the same position or a better one. Suppose Player A takes the whole cake and gives Players B and C \$10 each. Player A's valuation of his cake is the same because he received a \$30 cake and only gave away \$20. However, Player B previously only valued their slice at \$8, but now has \$10. The same is true for Player C, whose portion has gone from \$6 to \$10.

When the object being divided is not homogeneous, there are some cases where Pareto efficiency and envy-freeness are incompatible. In these cases, the division where everyone has the best possible value is not necessarily envy-free. For example, say there are 6 distinct objects and the three players value them in this order:

Player A: $3 \succ 1 \succ 6 \succ 2 \succ 5 \succ 4$
Player B: $2 \succ 6 \succ 1 \succ 3 \succ 5 \succ 4$
Player C: $5 \succ 3 \succ 1 \succ 4 \succ 6 \succ 2$

The envy-free division would consist of Player A with 3 and 6, Player B with 2 and 1, and Player C having 5 and 4. However, this envy-free division is not pareto-optimal, because if Player A instead took objects 3 and 1, and Player B took 2 and 6, then both players are better off with Player C's position unchanged. This new division is no longer envy-free because Player C can't be proven to value his division over Player A's.

4 Creating Polytopes of Slices

Polytopes can be used to represent all possible slicings of c cakes into s slices [2, 4]. Again, suppose the length of each cake is 1, and the sum of all slices in a single cake must be 1. These slices can first be arranged in a matrix with dimensions $(c \times s)$, with each row representing a different cake and each column representing a different slice. For example, x_{ij} would be the j-th slice of cake i.

$$\begin{bmatrix} x_{11} & x_{12} & x_{13} & \cdots & x_{1s} \\ x_{21} & x_{22} & x_{23} & \cdots & x_{2s} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x_{c1} & x_{c2} & x_{c3} & \cdots & x_{cs} \end{bmatrix}$$

A *pure slicing* is defined as a single slice that contains the whole cake. In our matrix, this would be shown with one slice of the cake being a 1, with every other slice as a 0. Any possible slicing of a cake can also be represented as a convex combination of its pure slicings. If these pure slicings were expressed in a polytope as vertices, then all slicings can be thought of as points within this polytope. We would only need to show the sizes of $s-1$ slices, because the last slice is already determined by the previous ones. Therefore, for one cake the number of dimensions needed is $s-1$. This number of dimensions is needed for every cake, so the dimension of the whole polytope is $c(s-1)$.

As stated earlier, each vertex of the polytope represents a pure slicing. Because there are s pure slicings on each cake, the total number of combinations of pure slicings is s^c, which is also the number of vertices of our polytope.

Now, let's consider the polytope of slicings for 2 cakes, each cut into 2 slices. The number of vertices is 2^2, or 4, and the dimension is $2(1) = 2$. The matrix would be 2×2 and a slicing which divides each cake in half would be $\begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}$. Another slicing which cuts one cake into slices of $\frac{1}{3}$ and $\frac{2}{3}$, and the second cake into $\frac{4}{5}$, $\frac{1}{5}$ would look like $\begin{bmatrix} \frac{1}{3} & \frac{2}{3} \\ \frac{4}{5} & \frac{1}{5} \end{bmatrix}$. Now, we can label these two slicings on the polytope.

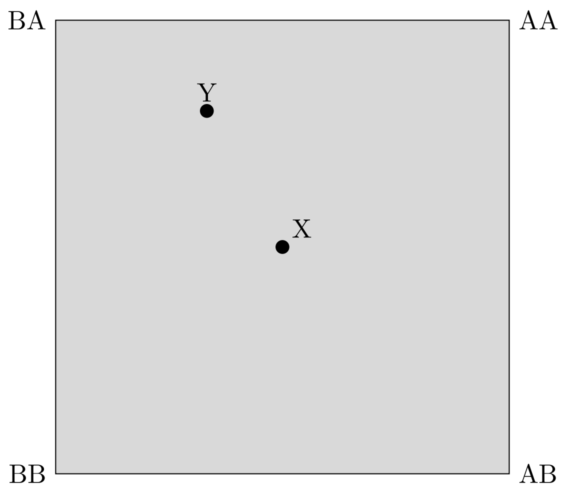

In the figure, point X represents $\begin{bmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{bmatrix}$ and point Y represents $\begin{bmatrix} \frac{1}{3} & \frac{2}{3} \\ 0.8 & 0.2 \end{bmatrix}$.

Each of the vertices represents a pure slicing, and each point on the square also represents some unique slicing. The vertices are also labeled to show which pure division is shown. The lower left corner is labeled BB, because both of the first slices are 0, meaning that for both cakes, the full cake is in slice B.

Now, let's label the polytope slightly differently. Instead of labeling corners with pure slices, they will instead be used to represent the slice selections of the players, labeled *bb*, *aa*, *ba*, *ab*. The divisions are still the same at the points in the square, and the corners show how each player chooses the slice with the full cake, which is why the letters have not changed. We will represent the preferences of players as coverings of the square in 4 pieces. Each covering contains one of the vertices, and every point in a covering represents a division where the player would choose the selection of that vertex. Here is an example of a player's covering for two homogeneous cakes.

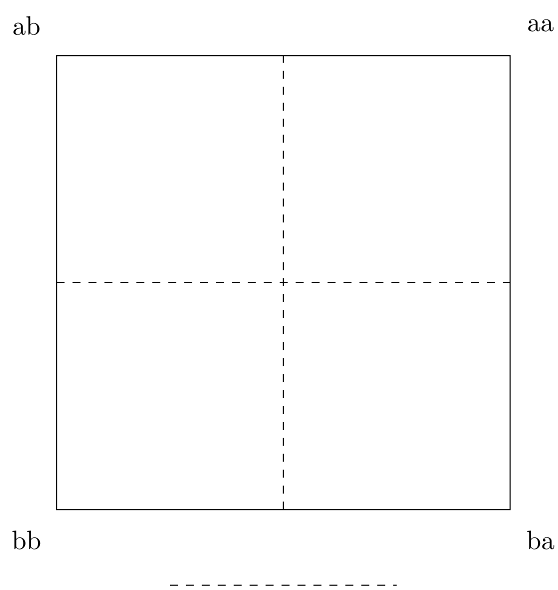

Player A's Preferences

In each of the smaller squares, the player chooses the slices shown in the corner associated with the square. For example, any division represented in the upper left square will lead to the player selecting slice a of cake 1 and slice b of cake 2. Notice how the player always chooses the bigger slice of each cake. With two players, as long as their coverings for opposite corners intersect, they can agree upon a slicing, so if both players have the covering shown, then the intersection would be in the center of the square, where there are two options for fair divisions.

Now, we can create a situation that does not involve homogenous cakes. Alice and Bob are both working for two days. Each day is split into a morning and evening shift. However, the end of the morning shift and beginning of the evening shift still must be determined for each day. Alice heavily prefers having two of the same shift, while Bob would rather have one of each. Each player would also choose a shorter shift over a longer one. We can create a covering for preferences in this problem, but first, the labels must be changed. The players prefer a smaller piece, or shorter shift, over a larger one. The labelings of the corners, a and b are then swapped. Here is a hypothetical covering for Alice and Bob.

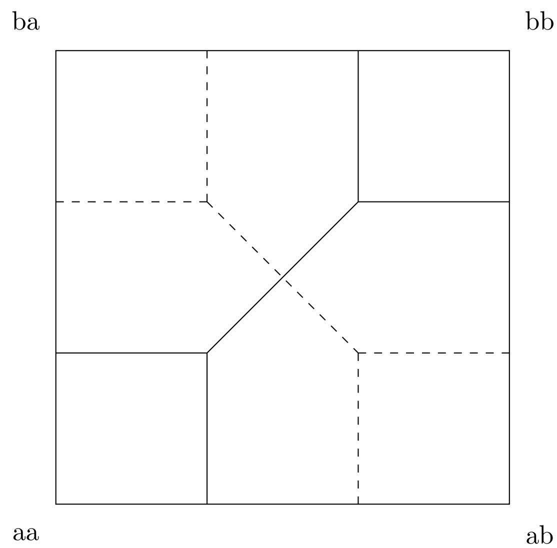

Player A's Preferences

Player B's Preferences

In this covering, Alice is willing to work a little longer to get two of the same shift, and is only willing to do two different shifts if they are both very short. Bob, on the other hand, would take longer shifts if it meant he would have two different ones, and he would only take two of the same shift if they were both short. For these preferences, there is no intersection of the two coverings that involve opposite corners. In other words, there is no division of the shifts on the two days where Alice and Bob each choose a different shift on both days. This means that given multiple cakes with multiple slices, there exists preferences of the players such that no slicing of the cakes allows for envy-free selections.

5 Sperner's Lemma

A *k-simplex* is a convex hull of $k+1$ affinely independent vertices. This would be a triangle generalized into k dimensions, beginning with a vertex, then a line segment, a triangle, tetrahedron, and then a 4-simplex, which would consist of 5 vertices, with each side being a tetrahedron.

A *triangulation* of a polytope P is any group of simplices whose union is P, and any intersection between two simplices is either empty or a shared face. A polytope is a geometric object with a finite number of sides sides. A polytope in two dimensions would be a polygon.

From a *triangulation* T of P, a *Sperner labeling* can be created under the following conditions:

- Each vertex of P is given a different label.
- Any vertex of T which lies on a facet or side of P is given the same label as one of the vertices of P that creates the facet. [2–4]

A *full cell* is any simplex of the same dimension as P in the triangulation T which has distinct labels for each of its vertices.

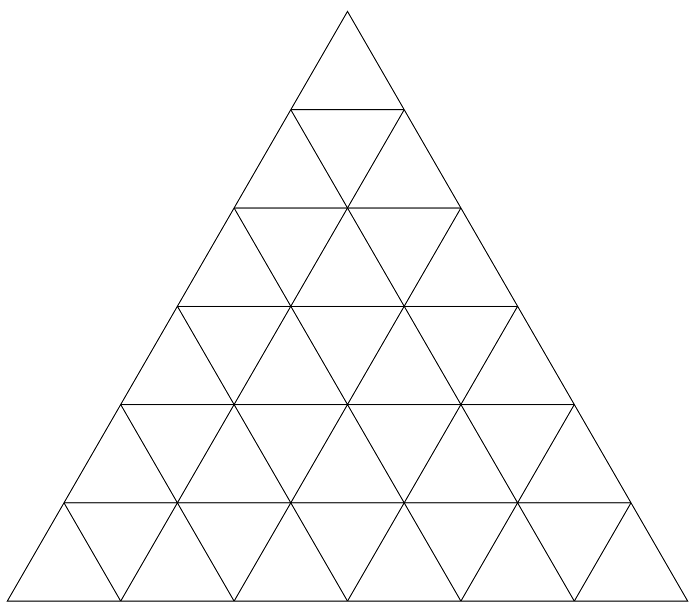

One example of a triangulation of an equilateral triangle

Lemma 1 (Sperner's Lemma)**.** *For any triangulation of a polytope P with a Sperner labeling, there exists an odd and non-zero number of full cells/*

Given a triangulation T of a triangle, such as the figure above, create a *Sperner labeling* with the 3 corners labeled with colors red, green, and blue.

The edges of T which are formed by blue and green can be treated as "doors" while all other edges are walls. Because the blue to green edge of the triangle begins with one color and ends with another, colors must alternate an odd number of times between the two vertices. This then means that the number of doors which lie on that edge must be odd, because any change in color would signify an edge in the triangulation T which starts with one color and ends with another. If a path is traced starting outside the triangle and only traveling through doors, then it will either hit a dead end somewhere within T, or leave the triangle through a door different from the one it started in. Because the number of outer doors is odd, there must be at least one path which does not lead back outside of the triangulation, and thus reaches a dead end, meaning it has entered a full cell.

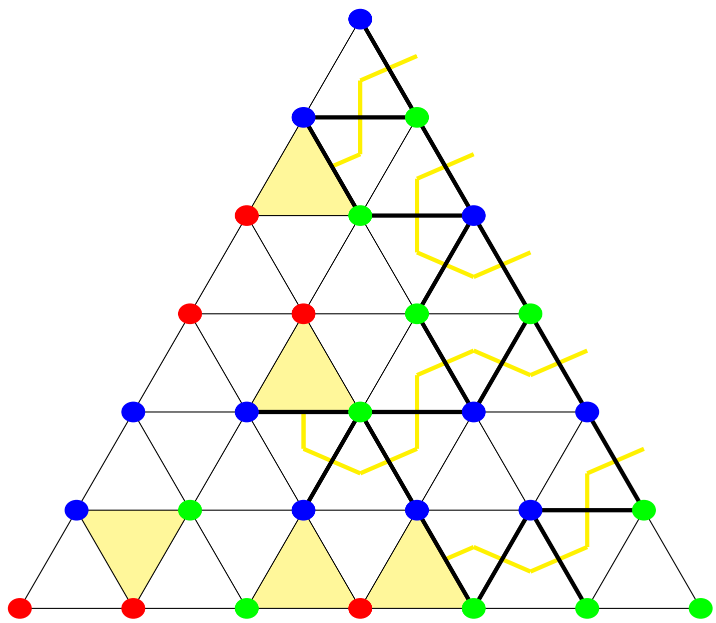

A *Sperner labeling* with paths highlighted

There is another proof of Sperner's Lemma that involves counting the total number of doors. We can use four values to count them:

- ℓ = Triangles with 1 door
- k = Triangles with 2 doors
- x = Doors that lie on the edge of P
- y = Doors that lie inside P

We can count the doors by adding the triangles with one door and double the triangles with two doors. However, this overcounts every door which is a part of two triangles. This can be accomodated for by subtracting y, as a door is part of two triangles if and only if it lies inside of P. This must then be equal to the total number of doors, or $x + y$.

$$\ell + 2k - y = x + y \tag{1}$$
$$\ell = x + 2y - 2k \tag{2}$$
$$\ell = 2(y - k) + x \tag{3}$$

Because we proved that x must be odd, the number of triangles with one door must be odd as well. Any triangle with only one door is a full cell, so this then means that the number of full cells is odd.

Now, imagine the problem of Alice, Bob, and Carl renting an apartment with 3 bedrooms, which they all value differently. How can they split the total rent and assign rooms among themselves so that for that division, no person believes that another room is better than their assigned one?

In the case of splitting rent, an equilateral triangle can be created using 3-dimensional coordinates using the equation $x + y + z = 1$, where x, y, and z are all positive. For any point (x, y, z) on this triangle, the values of x, y, and z represent the portion of the total rent paid by rooms 1, 2, and 3 respectively. A triangulation can be created with each of its vertices representing a particular different division of the rent among the rooms. Given each division, Alice and Bob can choose their preferred room, and these vertices can be labeled with two numbers (a, b), a being Alice's preference and b being Bob's. Each label can then be replaced according to the following assignment:

- $(1, 1)$, $(1, 2)$, or $(2, 1)$ becomes 3
- $(2, 2)$, $(2, 3)$, or $(3, 2)$ becomes 1
- $(3, 3)$, $(3, 1)$, or $(1, 3)$ becomes 2

Every vertex is now labeled with the room that neither Alice nor Bob wants. If it is assumed that all players prefer a free room over any other, then the triangulation now has a *Sperner labeling*. Sperner's Lemma says that there must be a simplex in this triangulation which has vertices of 1, 2, and 3, and by increasing the number of simplices, this full cell can be made infinitely small.

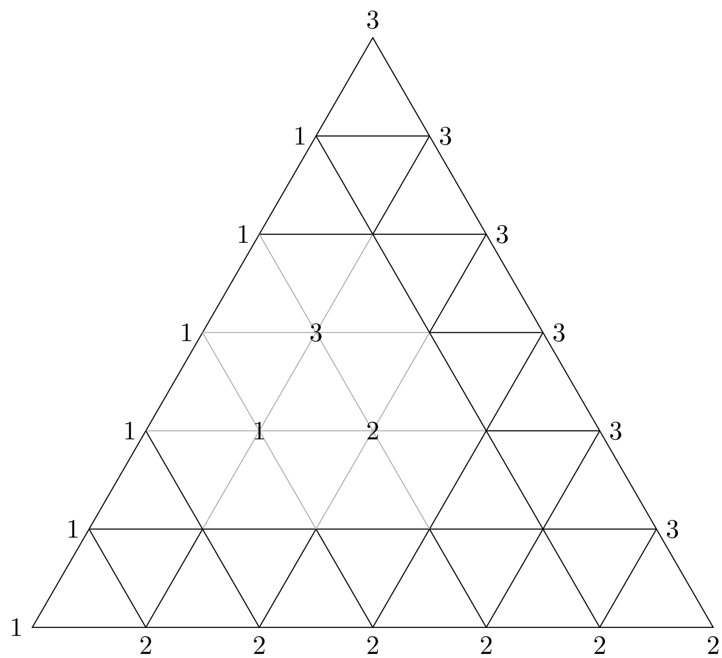

This full cell signifies that at each point, Alice and Bob leave a different room unchosen, meaning that Carl can choose any of the three rooms. Suppose Carl chooses room 3, but at the vertex labeled 3, Alice and Bob both preferred room 1. At the vertex labeled 1, someone must have preferred room 2, and because these vertices are so close together, these rent divisions are almost identical. The resulting division of rent makes it possible for each person to choose a different room they value as equal or greater than the other two.

6 KKM Lemma

Lemma 2 (KKM Lemma)**.** *Let Δ^n be an n-dimensional simplex with vertices $v_0, v_1, \ldots, v_n$. If there are $n+1$ closed sets $C_0, C_1, \ldots, C_n$ in Δ^n such that for each i, the set C_i contains all points of the simplex that are convex combinations of the vertices $\{v_0, v_1, \ldots, v_n\}$ excluding v_i, then the intersection of these sets $C_0 \cap C_1 \cap \cdots \cap C_n$ is non-empty.*

The KKM(Knaster–Kuratowski–Mazurkiewicz) Lemma [1, 6] can also be used in fair division problems in similar ways to both Sperner's Lemma and Polytope coverings. This Lemma uses coverings of simplices, stating that there must be some point where these three coverings intersect.

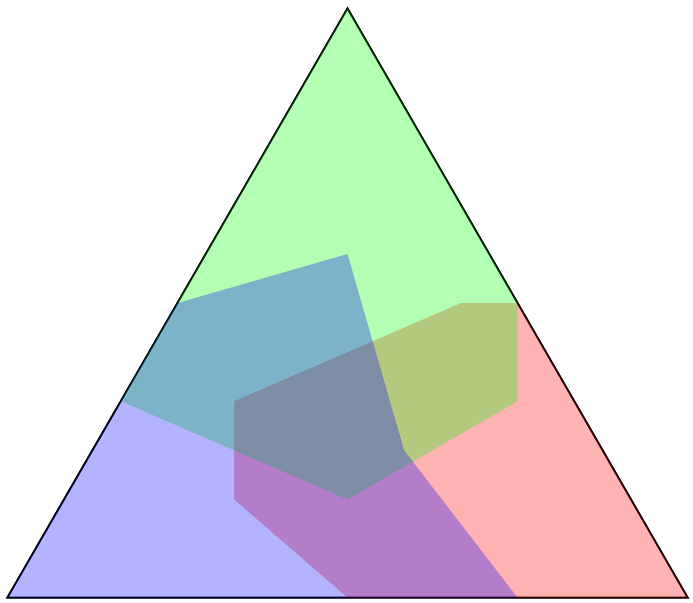

A *KKM Covering* of a triangle

Proof. Let $C_1, C_2, \ldots, C_{n+1}$ be closed subsets of the n-dimensional simplex, which we refer to as S_n. This means that each C_i is a distinct area within this shape. In this context, we want to understand how points in S_n relate to these subsets.

We define a function called γ. This function takes a point v from the simplex S_n and assigns it a number from the set $\{1, 2, \ldots, n+1\}$. For each point v, $\gamma(v)$ is the smallest number i such that i belongs to the face $F_E(v)$ of the simplex and the point v belongs to the set C_i. This means that $\gamma(v)$ helps us identify which subset contains the point while also considering the specific face of the simplex the point is associated with.

The function γ ensures that for every point v in the simplex S_n, the assigned number $\gamma(v)$ is part of the face $F_E(v)$. Additionally, if $\gamma(v) = i$, it confirms that the point v is included in the subset C_i.

Next, we consider a series of subdivisions of the simplex, denoted as $D_1, D_2, \ldots$. These subdivisions break down the simplex into smaller pieces, and we ensure that as we create these subdivisions, the size of the simplices within D_i gets smaller and smaller, with the diameter approaching zero.

In this process, we color every vertex v in each subdivision D_i using the value assigned by our function $\gamma(v)$. This creates a Sperner Labeling of the simplex, because if a vertex v lies in the I-facet (a specific face of the simplex), then its color, $\gamma(v)$, will belong to the face $F_E(v)$, which is a subset of I. This property ensures that we have a valid coloring scheme.

According to Sperner's Lemma, because of our Sperner Labeling, in each subdivision D_i, there must exist a simplex $S(V_i)$ such that the set of colors assigned to its vertices, $\gamma(V_i)$, contains every number from 1 to $n+1$. This means that this simplex represents all the subsets we are interested in, and as we continue our subdivisions, the size of the simplex $S(V_i)$ becomes smaller and smaller, approaching zero.

By the properties of the function γ, we can observe that for every index i (from 1 to n) and every index j in the set $\{1, 2, \ldots, n+1\}$, the intersection of the simplex $S(V_i)$ and the closed subset C_j is not empty. This indicates that there is at least one point that exists in both the simplex and the subset.

Let's introduce a new point, u_i, which is the average of all the vertices in the simplex V_i. This average point u_i is part of the simplex $S(V_i)$ and belongs to the n-dimensional simplex S_n as well.

Since the n-dimensional simplex S_n is bounded and closed, we can conclude that the sequence of average points u_i has a subsequence that converges to a limit point L. This limit point L is also located within the simplex S_n.

Now, every closed subset C_j has a property that for any small positive distance ϵ, there is at least one point within the set C_j that is within that distance from the limit point L. This implies that the limit point L must be inside the closed set C_j.

Putting this all together, we find that the limit point L belongs to the intersection of all the subsets $C_1, C_2, \ldots, C_{n+1}$. □

When visualized, the KKM Lemma looks very similar to Sperner's Lemma, and can be applied in similar ways. However, the KKM Lemma is also used to prove the Colorful KKM theorem, or Gale's theorem [1,6], which is very helpful in solving many types of fair division problems.

Theorem 1 (Gale's Theorem)**.** *Let $C_1, C_2, \ldots, C_{n+1}$ be closed subsets of the n-dimensional simplex S_n. If the sets satisfy the KKM condition, then there exists a point $x \in S_n$ that belongs to the intersection of the sets $C_{1,\pi_i(1)} \cap C_{2,\pi_i(2)} \cap \cdots \cap C_{d,\pi_i(d)}$ for all $i \in \{0, \ldots, d\}$, where π_i is a bijection mapping $\{1, \ldots, d\}$ to $\{0, \ldots, d\} \setminus \{i\}$.*

The Colorful KKM Theorem states that for any 3 KKM coverings of triangles, they can be overlaid on top of each other so that there is a single point that contains a different color on each covering.

Solutions to several fair division problems can be found using Gale's Theorem as well by representing divisions in a simplex, and then using coverings to represent a player's preferences in such a way that it creates a KKM covering.

We can use Gale's theorem to solve another fair division problem. Suppose Alice, Bob, and Carl are again deciding how to split work shifts, except now we are only focusing on a single day. There are three shifts, morning, evening, and night, and we must decide how to divide up the day into three time periods, so that each player prefers a different shift.

An equilateral triangle can be created again, with each point representing a different division. Unlike Sperner's Lemma, with Gale's theorem we can have each player create a different covering, with each color representing their preference at that division. Each edge again indicates that one of the shifts has a length of 0, so the players are guaranteed to select that shift at each of these points. Because the KKM lemma allows for intersections, we can use them to represent divisions where a player may not have a preference, and instead values two of the shifts equally.

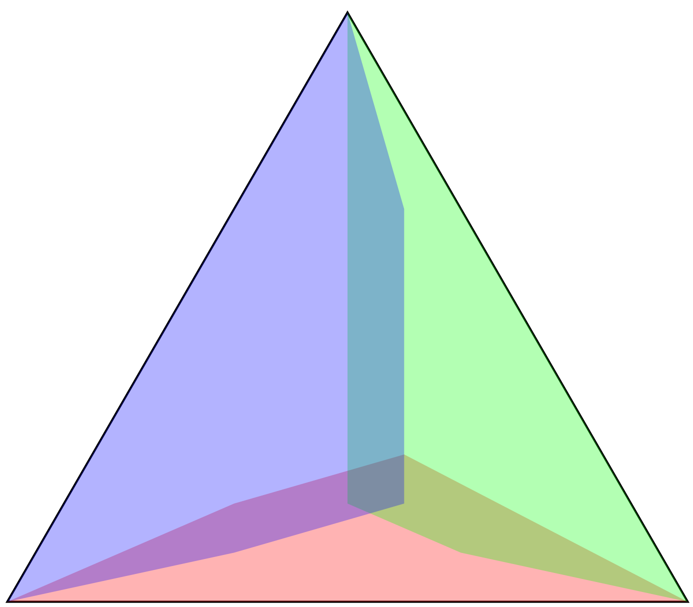

A possible covering for a player

In the figure above, the player prefers the shift represented by blue, which is shown by the size of the covering. By having each player create a different covering, Gale's theorem says that there must be some division where each player chose a different shift.

Notice that this theorem only applies to divisions of a single cake, and does not solve the previous work shift problem. However, Gale's theorem can be applied to any cake-cutting problem where each slice is distinct and may be valued differently to each player.

Another benefit of coverings instead of triangulations is that they can be mapped and rearranged to fit different situations. A covering which may not follow all the rules can be remapped so that it does. For example, here is a triangle with preference coverings that has 3 coverings on a single side.

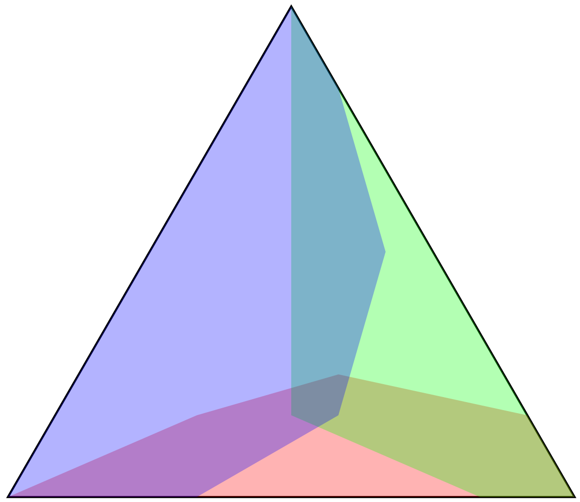

In this hypothetical covering, even though the preferences may not be realistic, it can still be rearranged into a KKM covering. By pushing the corners inward and pulling out the centers of each side, we can create a new triangle that is covered properly.

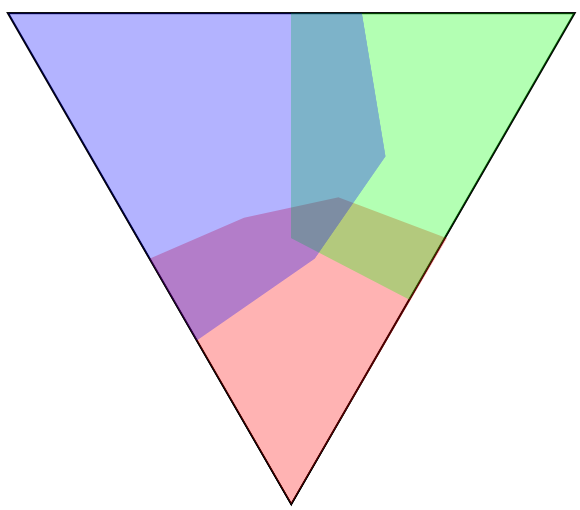

Like Sperner's Lemma, the KKM Lemma and Gale's theorem do not only apply to two dimensional simplices, and are generalized into n dimensions so they can be used for problems involving $n + 1$ players.

One issue with using such methods for divisions in real-world applications is that it sometimes may be impossible for players to give their preference for so many possible divisions. For example, Sperner's Lemma requires players to show their preference for hundreds of possibilities in order to produce an accurate envy-free divisions. On the other hand, KKM coverings may be easier for players to create, but it can still be challenging to picture all the divisions and decide what to choose while drawing out coverings. One solution to Sperner's Lemma is creating a computer algorithm that is able to find full cells without filling out the whole triangulation, which involves creating a path through doors as you go. However, many fair division solutions still do not have such simple real-world applications and truly are impossible to replicate with a finite amount of time.

7 Conclusion

In this paper, we have explored various fair division algorithms and topological methods to find envy-free divisions, demonstrating the complexities and intricacies that can be involved in achieving envy-free divisions among multiple players with different entitlements.

Beyond cake-cutting, fair division principles can be applied to numerous other domains. An interesting area for future exploration is the fair division of indivisible goods, where traditional methods of envy-free duvusuib may fall short. Meanwhile, there are whole other types of fair division problems which may not seek an envy-free division, but more optimal or efficient solutions. There are also problems of necklace-splitting which can be solved using various topological theorems.

Fair division is not simply a math problem to solve. It is used all the time, and when applied can greatly improve the efficiency of many people's lives. As the field of fair division evolves, plenty of interesting and practical problems remain to address. Extending our understanding of algorithms and their applications allows us to develop more robust and equitable solutions for a wide range of division challenges.

References

[1] Megumi Asada, Florian Frick, Vivek Pisharody, Maxwell Polevy, David Stoner, Ling Hei Tsang, and Zoe Wellner. Fair division and generalizations of sperner- and kkm-type results, 2017. https://arxiv.org/abs/1701.04955.

[2] John Cloutier, Kathryn L. Nyman, and Francis Edward Su. Two-player envy-free multi-cake division, 2009. https://arxiv.org/abs/0909.0301.

[3] Kiran Kedlaya. Fixed points and fair divisions, 2001. https://mathcircle.berkeley.edu/sites/default/files/BMC4/Handouts/bamc-cake.pdf.

[4] Justin J. Shin. *Envy-Free Fair Division With Two Players and Multiple Cakes*. Senior project, Bard College, 2016. https://digitalcommons.bard.edu/senproj_f2016/37.

[5] Seth Unruh. Envy-free divisions. *Rose-Hulman Undergraduate Mathematics Journal*, 10(2):Article 11, 2009. https://scholar.rose-hulman.edu/rhumj/vol10/iss2/11.

[6] Shira Zerbib. Topological methods in combinatorics. https://faculty.sites.iastate.edu/zerbib/files/inline-files/Topological_Methods_in_Combinatorics%20%2816%29.pdf.