FIXED POINT THEOREMS AND THEIR APPLICATIONS

TAE KYU KIM

1. Introduction

Fixed point theorems are powerful mathematical tools with a surprisingly wide range of applications. They tell us that for any sufficiently nice space X and any nice function $f : X \to X$, f has a fixed point i.e., a point $x^* \in X$ such that $f(x^*) = x^*$.

An example of a fixed point theorem is Brouwer's fixed point theorem (see Theorem 5.2). In this theorem, our space X is nice if it is homeomorphic to the closed unit ball in $\mathbb{R}^n$ for some $n > 0$, and the function f is nice if it is a continuous map. The philosophy behind fixed point theorems can be summarized with the following statement:

$$\text{Weaker hypothesis} \implies \text{harder, more general result.}$$

We haven't defined specifically what the space X or map f has to look like, only that they are sufficiently well-behaved. In general, the properties that we require are weak enough so that the fixed point theorems can be applied to many problems if we are clever in how we define the space X and the map f.

The Borsuk-Ulam theorem is a great example of this philosophy, though it isn't strictly a fixed point theorem. Although it is a general topological result, it has surprising (more concrete) consequences and applications to problems such as

- the avocado sandwich theorem,
- the Lyusternik-Shnirel'man theorem,
- and the necklace splitting problem.

In this paper, we state and prove two major fixed point theorems and showcase their applications in game theory. In particular, we use Sperner's lemma to prove that there is an envy-free division of a cake into n pieces for n people. Then, we use Brouwer's fixed point theorem to prove the existence of Nash equilibrium for non-cooperative games with finitely many players, each with finitely many pure strategies.

Much of the paper is dedicated to building up the technical definitions from simplicial geometry and game theory. However, this is a necessity, as understanding the application of fixed point theorems requires us to precisely define the space X and the map f.

2. Some Notions from Analysis

First, we review some results from analysis. There are many equivalent definitions for continuous functions. The most helpful definition for us will tell us about the convergence of sequences.

Definition 2.1. If $D \subseteq \mathbb{R}^n$ and $f : D \to \mathbb{R}^m$, then f is *continuous* if for every sequence $(x_j)_{j\geq 1}$ in D converging to $x \in D$, the sequence $(f(x_j))_{j\geq 1}$ converges to $f(x)$.

Continuous functions work well with inequalities:

Date: August 17, 2020.

Theorem 2.2. *If $f : D \to \mathbb{R}$ is continuous, $(x_j)_{j\geq 1} \to x$ in D, $(a_j)_{j\geq 1} \to a$ in $\mathbb{R}$ and $f(x_j) \leq a_j$ for all j, then $f(x) \leq a$.*

3. Sperner's Lemma

3.1. **Definitions from Simplicial Geometry.** In two dimensions, Sperner's lemma is a combinatorial fixed point theorem on vertex colorings of triangulations. A *triangulation* of a triangle Δ is a collection of triangles Δ_i such that

- the union of Δ_i is Δ,
- and if the intersection of Δ_i and Δ_j for $i \neq j$ is nonempty, then they intersect at a single vertex or a edge. If they intersect at an edge, its endpoints must be vertices of both Δ_i and Δ_j.

To discuss the n-dimensional version of Sperner's lemma, we must generalize the notion of triangles and edges to higher dimensions. The right notion of a n-dimensional triangle is an *n-simplex.*

Definition 3.1. An *n-simplex* is the convex hull in $\mathbb{R}^m$, with $m \geq n+1$, of $n+1$ geometrically independent points $v_0, v_1, \ldots, v_n$. We call $v_0, v_1, \ldots, v_n$ the *vertices of the simplex.*

The points $v_0, v_1, \ldots, v_n \in \mathbb{R}^m$ are geometrically independent if the vectors $\overrightarrow{v_0v_1}, \overrightarrow{v_0v_2}, \ldots, \overrightarrow{v_0v_n}$ are linearly independent.

Example. A 1-simplex is a line segment of some positive length. We require that it lies in $\mathbb{R}^m$ for some integer $m \geq 2$. A 2-simplex is a triangle with positive area. It must lie in $\mathbb{R}^m$ for some integer $m \geq 3$. A 3-simplex is a tetrahedron with positive volume, lying in $\mathbb{R}^m$ for some $m \geq 4$.

Nonexample. A square in $\mathbb{R}^4$ is almost a 3-simplex (it is the convex hull of $n+1$ points in $\mathbb{R}^m$ with $m \geq n+1$), but its vertices are not geometrically independent.

It will be useful to have a canonical form for an n-simplex.

Definition 3.2. The *standard n-simplex* in $\mathbb{R}^{n+1}$ is the convex hull of the points $(1, 0, \ldots, 0)$, $(0, 1, \ldots, 0), \ldots, (0, 0, \ldots, 1)$.

The reader should convince themselves that the standard 1-simplex is a line segment of length $\sqrt{2}$, and that the standard 2-simplex is an equilateral triangle of side length $\sqrt{2}$. The definition of convex hull implies that the standard n-simplex consists of all points $x = (x_0, x_1, \ldots, x_n) \in \mathbb{R}^{n+1}$ with $0 \leq x_i \leq 1$ for all $0 \leq i \leq n$ and $\sum_{i=0}^{n} x_i = 1$. More generally, if S is an n-simplex with vertices $v_0, v_1, \ldots, v_n$, then any point $x \in S$ can be uniquely expressed as

$$x = \alpha_0 v_0 + \alpha_1 v_1 + \cdots + \alpha_n v_n$$

for $\alpha_i \in \mathbb{R}$, $0 \leq \alpha_i \leq 1$ for all $0 \leq i \leq n$, and $\sum_{i=0}^{n} \alpha_i = 1$.

Suppose a triangle has vertices v_0, v_1, v_2. The triangle contains lower-dimensional components called edges and vertices. Notice that edges are formed by taking the 1-simplices with vertex sets $\{v_0, v_1\}$, $\{v_0, v_2\}$, or $\{v_1, v_2\}$. The vertices are 0-simplices with vertices at v_0, v_1, or v_2. We may similarly define the lower-dimensional components of an n-simplex by taking the convex hull of a subset of the vertex set.

Definition 3.3. An *m-face* of an n-simplex S is the m-simplex formed by $m+1$ vertices out of the $n+1$ vertices of S. An $(n-1)$-face of an n-simplex is called a *facet.*

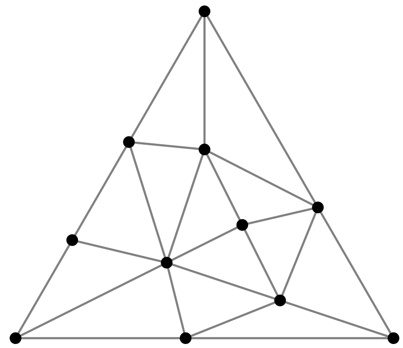

Figure 1. Subdivision of a 2-simplex i.e., a triangulation.

Exercise. For given m, n, how many m-faces does an n-simplex have?

Now, we define subdivisions of n-simplices, which are a generalization of triangulations for higher dimensions. Just like in the two-dimensional case, we want the sub-simplices to have the same dimensions as the original simplex (a triangle is subdivided into triangles; an n-simplex is subdivided into n-simplices).

Definition 3.4. A *subdivision* of an n-simplex S is a collection of subsets of S, each an n-simplex, called a *sub-simplex* of S, such that:

(1) The union of all the sub-simplices is S.
(2) Any two sub-simplices either do not intersect or have an intersection that is a common face.

We call a facet of a sub-simplex a *sub-facet.* A subdivision of a 2-simplex is called a *triangulation.*

See Figure 1 for an example triangulation.

3.2. **Sperner's Lemma.** With the preliminary definitions out of the way, we are ready to state and prove Sperner's lemma. In the two-dimensional case, each vertex was labeled a color from $\{0, 1, 2\}$. In the n-dimensional case, each vertex is labeled with a color from $\{0, 1, \ldots, n\}$.

It will be helpful to categorize certain types of simplices by the colors of their vertices. For analyzing the subdivision of an n-simplex, we will only care about describing sub-simplices and sub-facets in terms of their vertices' colors.

Definition 3.5. Assume we have a subdivision of an n-simplex S and that the vertices of sub-simplices of S are labeled with colors from $\{0, 1, \ldots, n\}$. A facet in S is called a $(a_0, a_1, \ldots, a_{n-1})$ *facet* if its vertices are labeled $a_0, a_1, \ldots, a_{n-1}$. Also an n-simplex in S (including S and sub-simplices) is called an $(a_0, a_1, \ldots, a_n)$ *simplex* if its vertices are labeled $a_0, a_1, \ldots, a_n$.

Example. Suppose S were a $(0, 1, 2, \ldots, n)$ simplex. Its $n+1$ facets have vertex colorings $(0, 1, 2, \ldots, \hat{k}, \ldots, n)$ where the hatted element k is excluded from the sequence.

Definition 3.6. Given a subdivision of an n-simplex S with the vertices of the sub-simplices labeled with colors from the set $\{0, 1, \ldots, n\}$, the subdivision is said to have a *Sperner labeling* if the sub-simplices are labeled according to the following rules:

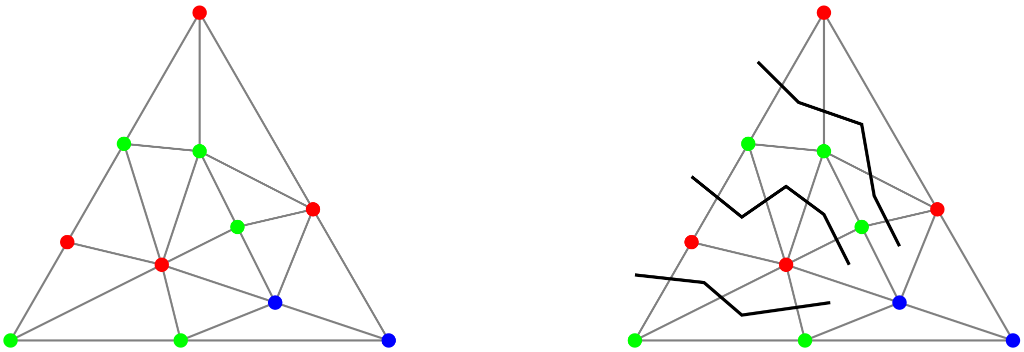

(a) Sperner labeling of the triangulation from Figure 1. Red = 0, Green = 1, Blue = 2.

(b) The unique complete collection of p-paths for the Sperner labeling.

Figure 2. A Sperner labeling and its corresponding p-paths.

(1) S is a $(0, 1, \ldots, n)$ simplex,
(2) The vertices of sub-simplices on a facet of S do not have the same label as the vertex opposite the facet.

Example. A valid Sperner labeling is given in Figure 2a. We can verify that S is a $(0, 1, 2)$ simplex and that the vertices on the edges of S do not have the same label as the vertex opposite that edge. For instance, the vertices on the left edge must be labeled 0 or 1 (respectively red or green) because the bottom-right vertex is labeled 2 (respectively blue).

Example. Suppose we subdivided a 3-simplex i.e., a tetrahedron. Some of the sub-simplices will have vertices on the faces (more precisely, the facets) of S. We claim that the subdivision induces triangulations of the faces of S. For a Sperner labeling, the 4 vertices of S must be labeled $0, 1, 2, 3$ in some order. Then, each face f of S is the convex hull of 3 of the 4 vertices of S. For sake of concreteness, let f be a $(0, 1, 2)$ facet. The vertices contained in f (from the induced triangulation) must be labeled with a 0, 1, or 2. They cannot be labeled 3 as the vertex of S opposite to f is labeled with a 3.

Remark 3.7. There's actually an equivalent formulation of Sperner labeling that may aid our intuition:

- all vertices of P have distinct labels,
- and the label of any vertex of a sub-simplex of P that lies on a facet of P, must match the label of one of the vertices of P that spans that facet.

The reader should convince themselves that the two definitions for a Sperner labeling are equivalent for subdivisions of a 2-simplex and 3-simplex. This alternative definition is how we actually generalize Sperner labeling for subdivisions of polyhedra.

We may consider a graph dual to the subdivision: define the vertices of the graph to be the centers of the sub-simplices, and let an edge connect two vertices if their corresponding sub-simplices intersect at a sub-facet. For the counting argument in the proof of Sperner's lemma, we consider certain paths on this dual graph (however, we also let the paths leave the n-simplex).

Definition 3.8. A *p-path* (permissible path) in an n-simplex is defined as a path that:

(1) begins in a $(0, 1, \ldots, n)$ sub-simplex or outside S,

(2) ends in a $(0, 1, \ldots, n)$ sub-simplex or outside S,
(3) crosses from one region to an adjacent region only through a $(0, 1, \ldots, n-1)$ sub-facet,
(4) crosses each $(0, 1, \ldots, n-1)$ sub-facet exactly once.

A collection of p-paths is *complete* if every $(0, 1, \ldots, n-1)$ sub-facet is crossed by a unique path in the collection.

Let's understand this definition for $n = 2$. Figure 2a shows a Sperner labeling of a triangulation. We connect the centers of two adjacent sub-triangles if their shared edge is colored $(0, 1)$. Figure 2b shows the end result of this process. Notice that we obtain a collection of paths which begin and end in a $(0, 1, 2)$ sub-simplex or outside S. Moreover, every $(0, 1)$ sub-facet is crossed exactly once by the paths. Thus, the collection of p-paths is complete.

For higher dimensions, we connect the centers of two adjacent sub-simplices if their shared sub-facet is colored $(0, 1, \ldots, n-1)$. In a similar fashion as before, this process will give us a complete collection of p-paths. Thus, we claim the following proposition.

Proposition 3.9. *A complete collection of p-path exists for any n-simplex.*

We leave the proof as an exercise to the reader. We are now ready to prove Sperner's lemma for n-simplices.

Theorem 3.10 (Sperner's lemma)**.** *In any subdivision of an n-simplex S that has a Sperner labeling ($n \in \mathbb{Z}_{>0}$), there exists an odd number of $(0, 1, \ldots, n)$ sub-simplices.*

Proof. We proceed by induction on n. For $n = 1$, we are given a line segment $[0, 1]$ and some points $a_0, \ldots, a_k \in \mathbb{R}$ with $0 = a_0 < a_1 < a_2 < \cdots < a_{k-1} < a_k = 1$. Each distinguished point including the endpoints is colored with a 0 or 1 such that the colors of the endpoints are different. We wish to prove that there are an odd number of $(0, 1)$ sub-simplices. This corresponds to finding the number of indices i such that a_i and a_{i+1} have different colors.

When $k = 1$, this is clear. We consider how the parity changes when we add a new point to the sequence. Say we insert a_i between a_{i-1} and a_{i+1}. We can check that no matter what the colors of a_i, a_{i-1}, a_{i+1} are, the parity does not change. Hence, there is an odd number of indices i such that a_i and a_{i+1} have different labels and so there are an odd number of $(0, 1)$ sub-simplices.

Suppose that Sperner's lemma holds for subdivisions of $(n-1)$-simplices. We will prove Sperner's lemma for n-simplices, so assume that we have a subdivision of an n-simplex S with Sperner labeling. By Remark 3.7, any $(0, 1, \ldots, n-1)$ sub-facet that lies on a facet of S must lie on the $(0, 1, \ldots, n-1)$ facet of S. The $(0, 1, \ldots, n-1)$ facet of S is an $(n-1)$-simplex with (an induced) Sperner labeling. Thus, by induction, the $(0, 1, \ldots, n-1)$ facet has an odd number of $(0, 1, \ldots, n-1)$ sub-facets. Hence, the number of $(0, 1, \ldots, n-1)$ sub-facets on the boundary of S is odd.

From Proposition 3.9, we have a complete collection of p-paths for the subdivision of S. Each path begins/ends inside a unique $(0, 1, \ldots, n)$ sub-simplex or outside S. Consider the number of $(0, 1, \ldots, n-1)$ sub-facets on the boundary of S that are crossed by any p-path. The only way that a p-path can pass through the boundary of S is through a $(0, 1, \ldots, n-1)$ sub-facets on the boundary of S, and each time the p-path passes through the boundary, the path switches between being inside and outside S. Thus, the p-path has endpoints both inside and outside S if and only if it passes through an odd number of $(0, 1, \ldots, n-1)$ sub-facets on the boundary of S. Call this an inside-outside path.

Any other type of paths will pass through the boundary an even number of times. Since the collection of p-paths passes through all $(0, 1, \ldots, n-1)$ sub-facets exactly once and we know that there is an odd number of $(0, 1, \ldots, n-1)$ sub-facets on the boundary of S, there must be an odd number of inside-outside paths. Finally, if a p-path ends inside S, then it ends in a $(0, 1, \ldots, n)$ sub-simplex. This implies that the number of $(0, 1, \ldots, n)$ sub-simplices is odd. ∎

4. Envy-Free Division

In this section, we will use Sperner's lemma to prove the existence of envy-free division of a single cake to n people. We model the cake as the line segment $[0, 1]$ and make $n - 1$ cuts to get pieces of sizes $x_1, x_2, \ldots, x_n$. Notice that $x_i \geq 0$ for all i and $\sum_{i=1}^{n} x_i = 1$. Thus, the space of possible cuts into n pieces forms an $(n-1)$-simplex. More formally,

Definition 4.1. A *cake-cut* for n people is defined as a point $(x_1, x_2, \ldots, x_n)$ where x_i is the size of the i^{th} piece, such that

$$0 \leq x_i \leq 1 \quad \forall i \in \{1, 2, \ldots, n\},$$

$$\sum_{i=1}^{n} x_i = 1.$$

The set of all possible cake-cuts S is a $(n-1)$-simplex.

For each particular cutting of the cake, each person P_α will have a preference among the n pieces. Moreover, a person may be impartial between multiple pieces, so they are fine with getting any one of their top choices. Maybe person P_1 only cares getting the biggest piece, or maybe person P_2 care about getting the piece with the most strawberries. Each person may have wildly different preferences, but we want to deal with a "preference" function that is somewhat nice to work with.

Definition 4.2. For $\alpha = P_1, P_2, \ldots, P_n$, and $x = (x_1, x_2, \ldots, x_n) \in S$, let $C_\alpha(x)$ be a subset of $\{1, 2, \ldots, n\}$. We call the set-valued function C_α a *choice function for person* α if:

1. $C_\alpha(x) \neq 0$ or all $x \in S$.
2. For all $x \in S$, if $x_i = 0$, then $i \notin C_\alpha(x)$.
3. C_α is continuous.

The first two restrictions make intuitive sense. The first condition says that each person will like at least one of the pieces. The second condition guarantees that a nonempty piece is always preferred over an empty piece. We can imagine that everyone is hungry, so they would rather eat a very small piece than eat nothing. The third condition is a bit more technical and will come into play in the proof of the next theorem.

Theorem 4.3 (Existence of Envy-Free Division)**.** *Let $C_{P_1}, C_{P_2}, \ldots, C_{P_n}$ be the choice functions on S. There exists $x \in S$ and distinct $\alpha_1, \alpha_2, \ldots, \alpha_n$ among $P_1, P_2, \ldots, P_n$ such that $i \in C_{\alpha_i}(x)$ for all $1 \leq i \leq n$. Thus, at $x \in S$, each person can have a piece of cake that they choose.*

In the statement of the above theorem, $(\alpha_1, \alpha_2, \ldots, \alpha_n)$ is a permutation of $(P_1, P_2, \ldots, P_n)$. This condition says that person i is allowed to take the j^{th} piece where $i \neq j$, as long everyone takes different pieces.

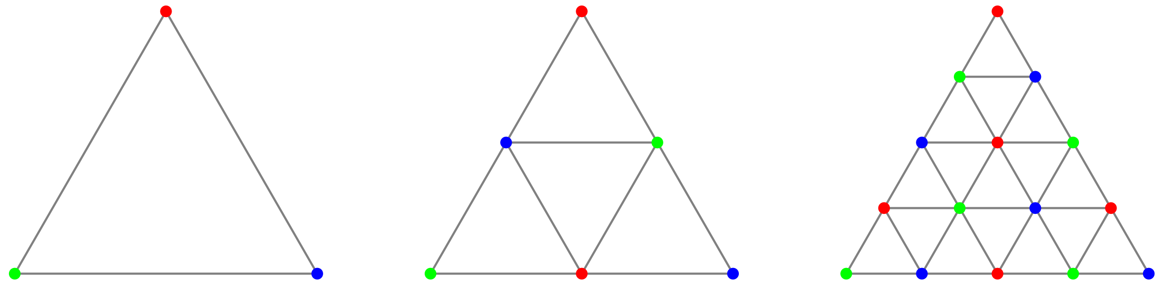

(a) $\mathcal{T}_1$ **(b)** $\mathcal{T}_2$ **(c)** $\mathcal{T}_3$

Figure 3. Increasingly finer triangulations and their Sperner labelings.

Proof. We start with a sequence of increasingly finer subdivisions $\mathcal{T}_1, \mathcal{T}_2, \ldots$ of S, constructed so that every sub-simplex of each subdivision is labeled $(P_1, \ldots, P_n)$. For $n = 2$, we can form a finer triangulation T_{i+1} by adding the midpoints of the sub-facets of T_i as vertices of the sub-simplices. The first two steps of this process are shown in Figure 3. A generalization of this subdivision algorithm for simplices is called the barycentric subdivision. A nice property is that the vertices of the resulting subdivision of an $(n-1)$-simplex can be labeled with n numbers so that every sub-simplex has vertices of every label. Now, for each vertex of the triangulation, we replace the label P_i with a label i such that $i \in C_{P_i}(x)$ where x is the coordinate of the vertex. If the choice function consists of more than one element, then choose one arbitrarily.

We claim that this new labeling is a Sperner labeling. The vertices of S represent cake cuts where the i^{th} piece is the entire cake and the other pieces are empty. So any person would prefer the i^{th} piece, giving the vertex label i. Hence, S is a $(1, \ldots, n)$ simplex. Now consider any vertex on a facet of S; for concreteness, we will let v be a vertex on the $(1, 2, \ldots, m-1)$ facet of S. The coordinate of v is a linear combination of the coordinates of the vertices of S labeled $1, 2, \ldots, m-1$. In particular, the m^{th} coordinate of the vertices of S that span over the facet is 0, so the m^{th} coordinate of v must also be 0. Thus, any person would never prefer the m^{th} piece for the cake cut represented by v. Thus, we have shown that the new labeling is a Sperner labeling.

By Sperner's lemma, there is at least one $(1, 2, \ldots, m)$ sub-simplex in any triangulation $\mathcal{T}_n$ of S. Denote this sub-simplex T_n. Any $(1, 2, \ldots, m)$ sub-simplex tells us that person P_α prefers the $\sigma(\alpha)^{\text{th}}$ piece where $\sigma \in S_n$ is a permutation. Hence, each T_n "belongs" to a unique permutation in S_n. Since there are finitely many permutations and infinitely many $(1, 2, \ldots, m)$ sub-simplices, there is at least one permutation σ for which infinitely many $(1, 2, \ldots, m)$ sub-simplices belong to. Without loss of generality, we will let σ be the identity permutation so that $\sigma(\alpha) = \alpha$ for all $\alpha = 1, 2, \ldots, n$. Moreover, for simplicity, we will let all of the T_n belong to the identity permutation.

Consider the sequence of vertices of T_n with label 1. Since this is an infinite sequence of points in a compact subset of $\mathbb{R}^n$, there exists a subsequence of points that converges to limit point in S. For simplicity, let the subsequence be $(x_n)_{n\geq 1}$ where x_n is a vertex of T_n, and let x be the limit of x_n. As $\mathcal{T}_n$ is a sequence of increasingly finer subdivisions, the diameters of T_n goes to 0 as n increases. Hence, all of the vertices of T_n converge to x. By construction, for every vertex $x_{\alpha,n} \in T_n$ with label α, we have $\alpha \in C_\alpha(x_{\alpha,n})$. Continuity of the choice function tells us that $\alpha \in C_\alpha(x)$ since $(x_{\alpha,n})_{n\geq 1}$ converges to x. Hence, for the cake cut represented by x, person α will prefer the α^{th} piece, as desired. ■

5. Brouwer's Fixed Point Theorem

In this section, we will state and prove the Brouwer's fixed point theorem, which is the topological analogue of Sperner's lemma.[1] First we define a closed unit ball in $\mathbb{R}^n$.

Definition 5.1. Let $\overline{B^n}$ denote the closed unit ball in $\mathbb{R}^n$. This is the set of points in $\mathbb{R}^n$ that at most distance 1 from the origin. Formally,

$$\overline{B^n} = \left\{ (x_1, x_2, \dots, x_n) \in \mathbb{R}^n : \sqrt{x_1^2 + x_2^2 + \cdots + x_n^2} \leq 1 \right\}.$$

Theorem 5.2 (Brouwer)**.** *For any continuous function $f : \overline{B^n} \to \overline{B^n}$, there is a point x^* such that $f(x^*) = x^*$.*

Proof. Because the standard n-simplex S is homeomorphic to $\overline{B^{n-1}}$, it suffices to prove that any continuous function $f : S \to S$ has a fixed point. Let $\mathcal{T}_1, \mathcal{T}_2, \dots$ be a sequence of increasingly finer subdivisions of S. Assume that the diameters of the sub-simplices tend to 0 as $i \to \infty$.

If any vertex of a sub-simplex in a subdivision is a fixed point of f, then we are done since f has a fixed point. So, assume that none of the vertices of the sub-simplices of the subdivisions are fixed points of f. For any $x = (x_0, x_1, \dots, x_n) \in S$, let

$$f(x) = f(x_0, x_1, \dots, x_n) = (f_0(x), f_1(x), \dots, f_n(x)).$$

Assuming that x is not a fixed point, we can label x with the number p such that

$$p = \min\{k = 0, 1, \dots, n | f_k(x) < x_k\}.$$

We claim that this labeling is a Sperner labeling on the subdivisions $\mathcal{T}_i$. First note that by assumption, the vertices of the sub-simplices are not fixed points, so the labels are well-defined. The vertices of S are $(1, 0, \dots, 0), (0, 1, 0, \dots, 0), \dots, (0, 0, \dots, 1)$, and according to the labeling, they must be labeled $0, 1, \dots, n$, respectively. Thus, S is a $(0, 1, \dots, n)$ simplex. Moreover, suppose that x is a point on the $(0, 1, \dots, \hat{k}, \dots, n)$ facet, which is opposite to the vertex of S with a 1 in the k^{th} coordinate. The k^{th} coordinate of the vertices that span the facet is 0, so k^{th} coordinate of x is also 0. Hence, x will not be labeled k. This proves that we have a Sperner labeling on every subdivision $\mathcal{T}_i$.

By Sperner's lemma, there exists at least one $(0, 1, \dots, n)$ sub-simplex in S for each $\mathcal{T}_i$. For each $\mathcal{T}_i$, choose a $(0, 1, \dots, n)$ sub-simplex, then let v_k^i be the vertex of the sub-simplex labeled k. We obtain $n + 1$ sequences of points in S: $(v_k^i)_{i \geq 1}$ for $k = 0, 1, \dots, n$. For now, consider just $k = 0$. Since S is a compact space, there exists a subsequence of $(v_0^i)_{i \geq 1}$ that converges to a point v^*. For simplicity, assume that the subsequence is the entire sequence $(v_0^i)_{i \geq 1}$.

Since the diameter of the sub-simplices in $\mathcal{T}_i$ goes to 0 as $i \to \infty$, all of the sequences $(v_k^i)_{i \geq 1}$ also converge to v^*. To prove that v^* is a fixed point, let $v^* = (v_0^*, v_1^*, \dots, v_n^*)$ and let $v_{k,m}^i$ denote the m^{th} coordinate of v_k^i. Note that for all k, m, $(v_{k,m}^i)_{i \geq 1}$ converges to v_m^*.

[1]Sperner's lemma, Brouwer's fixed point theorem, and the Knaster–Kuratowski–Mazurkiewicz lemma are three equivalent results from three different areas of math. Sperner's lemma is combinatorial, using a parity argument and admitting an induction proof. Proving Brouwer's fixed point theorem generally requires using Sperner's lemma or machinery from topology. The KKM lemma is a result on set coverings and can be proved from the other two variants.

Now, v_k^i has the labeling k, so by the labeling we must have $f_k(v_k^i) < v_{k,k}^i$. Since f is continuous, so is f_k and so by Theorem 2.2, the limit point v^* must satisfy

$$f_k(v^*) \leq v_k^*. \tag{5.1}$$

Notice that we no longer have strict inequality, but in fact we actually want to show equality. Since (5.1) holds for all k and $\sum_{k=0}^{n} f_k(v^*) = \sum_{k=0}^{n} v_k^* = 1$ (because the points lie in the standard n-simplex), we must have equality in (5.1) for all k. Hence, $f(v^*) = v^*$, so v^* is a fixed point of f, as desired. ■

6. Existence of Nash Equilibrium

In this section, we will prove the existence of Nash equilibrium using Brouwer's fixed point theorem. This is a major theorem in economics and game theory. After one knows that Nash equilibrium exist in certain games, it's natural to study the quality of the equilibrium points and ask whether those are reasonable outcomes. For instance, an equilibrium point may be unstable; a small perturbation may incentivize players to move to a new equilibrium point. Thus, while a reasonably simple result, the guaranteed existence of Nash equilibrium leads us to deeper questions in game theory.

Remark 6.1. In 1950, John Nash proved the existence of Nash equilibrium using the Kakutani fixed point theorem. In 1951, he proved it again using Brouwer's fixed point theorem. Since the Kakutani fixed point theorem is more general version of Brouwer's fixed point theorem, the second proof was an improvement over the first proof.

6.1. **Game Theory Background.** Nash equilibrium is John Nash's proposed solution to a non-cooperative game. In a non-cooperative game,

- each player knows other players' strategies,
- and each player wants to maximize their expected payoff.

A combination of the players' strategies is said to be a *Nash equilibrium* if no player can increase their expected payoff by changing their own strategy if the other players keep theirs unchanged. That is, no player has an incentive to change their strategy if everyone else keeps their strategies the same.

We will refer back to the following example game repeatedly throughout this section.

Example. Alice and Bob play the following game:

- Each person tosses in a nickel or quarter into a hat.
- If the two coins from Alice and Bob were of the same type, then Alice gets both coins. If they were of different value, then Bob gets both coins.

In general, each player prefers to toss in a nickel instead of a quarter, since playing a nickel has a lower risk for the same profit margin. However, the players' goals are in conflict with each other, so just through the rules of the game, Alice and Bob are adversaries. A non-cooperative game does not generally have to be adversarial, or even zero-sum; the quality of the game depends on the set of possible player actions and the payoff functions for the players.

With the intuitive explanation in mind, we will build up the technical definition for Nash equilibrium. First, we introduce pure and mixed strategies.

Definition 6.2. Each player i has a set of *pure strategies* $\{\pi_{i\alpha}\}$. Each pure strategy corresponds to doing a particular action with probability 1.

Example. A pure strategy for Alice in the example game would be playing a nickel with probability 1. Her other pure strategy (she has two) is playing a quarter with probability 1. However, this is a dumb strategy because Bob has a counter-strategy that always lets him win: play a quarter or a nickel with probability 1. In general, a pure strategy is not a good strategy because the other players can exploit the predictability of a pure strategy.

How can a player make sure that the other players can't predict his or her moves? One way is to assign probabilities to each of the pure strategies and play those strategies with that assigned probability.

Definition 6.3. A **mixed strategy** is given by a collection of non-negative numbers $\{c_{i\alpha}\}$ such that $\sum_\alpha c_{i\alpha} = 1$. Player i plays the pure strategy $\pi_{i\alpha}$ with probability $c_{i\alpha}$.

From this formulation, it is clear that the set of strategies for any player is a k-simplex for some k. For precisely,

Proposition 6.4. *Suppose player i has k pure strategies available to him. The set of strategies for player i is the standard $(k-1)$-simplex with k vertices*

$$\{(1,0,0,\ldots,0),(0,1,0,\ldots,0),\ldots,(0,0,\ldots,1)\}.$$

If we have finitely many players—say n players—then the set of all strategy combinations of the n players is the direct product of the $(k-1)$-simplices from each player. This direct product is homeomorphic to $\overline{B^m}$ for some integer $m > 0$.

Now that we understand the players' possible actions, let's see what their objectives are. Their goals are to maximize the value of their individual payout function:

Definition 6.5. The payoff is a function of all of the players' strategies:

$$p_i(\$) = p_i(s_1, s_2, \ldots, s_n)$$

is the expected payoff for player i if player j plays strategy s_j $(1 \leq j \leq n)$.

Example. We can easily understand the payouts for pure strategies for the example game. The following table tells us (payout to Alice, payout to Bob) for the combinations of pure strategies available to Alice and Bob.

		Bob	
		N	Q
Alice	N	$(5,-5)$	$(-5,5)$
	Q	$(-25,25)$	$(25,-25)$

While not obvious at first, we want the payout functions to be n-linear. That is, if $s_i = \sum_\alpha c_{i\alpha}\pi_{i\alpha}$, then for all payout functions p_j,

$$p_j(\$) = \sum_\alpha c_{i\alpha} p_j(s_1, \ldots, s_{i-1}, \pi_{i\alpha}, s_{i+1}, \ldots, s_n).$$

The idea is similar to the linearity of expected value. Each combination of pure strategies gives us a particular payout. Assigning probabilities to those pure strategies gives us probabilities for combinations of pure strategies. Hence, to calculate the expected payoff, we should multiply the payout for the combination of pure strategies by the probability that the particular combination of pure strategies will be played. The linearity of the payout functions implies that the payout function is uniquely determined by its values on the combinations of pure strategies.

6.2. **Nash equilibrium.** Now, we are ready to formally define Nash equilibrium.

Definition 6.6 (Nash equilibrium)**.** A *Nash equilibrium* is a strategy combination $\mathbb{S}$ such that for all i,

$$p_i(\mathbb{S}) = \max_{s \in S_i}\{p_i(s_1, s_2, \ldots, s_{i-1}, s, s_{i+1}, \ldots, s_n)\},$$

where S_i is the set of all strategies available to player i. That is, no player can increase their expected payoff by changing their choice, if the other players' strategies are fixed.

Theorem 6.7. *Any non-cooperative n-player game in which each player has finitely many pure strategies and any mixed strategy of the pure strategies exists, has a Nash equilibrium.*

Example. Before we prove the existence of Nash equilibrium for non-cooperative games, let's calculate the Nash equilibrium points of the example game. Suppose Alice plays a nickel with some probability $p \in [0, 1]$. With what probability q should Bob play a nickel so that the game is in Nash equilibrium? The payout functions for Bob when $q = 0, 1$ are

$$p_B((p, 1-p), (0, 1)) = 5p - 25(1-p), \quad p_B((p, 1-p), (1, 0)) = -5p + 25(1-p).$$

If Bob plays a nickel with probability q, then the corresponding payout function is

$$p_B((p, 1-p), (q, 1-q)) = (1-q) \cdot p_B((p, 1-p), (0, 1)) + q \cdot p_B((p, 1-p), (1, 0)).$$

Notice that if the payouts for $q = 0, 1$ are different for a given value of p, then Bob's optimal strategy is picking $q = 0$ or $q = 1$, depending on which of the payout functions is greater. However, Alice's optimal strategy is then picking $p = 0$ or 1. Since Alice and Bob's goals conflict, picking pure strategies will never give us a Nash equilibrium.

Thus, we should look at when $p_B((p, 1-p), (0, 1)) = p_B((p, 1-p), (1, 0))$, so that any of Bob strategies are optimal. The payouts are equal when $p^* = \frac{5}{6}$. With a similar argument, we find that the Nash equilibrium should have $q^* = \frac{1}{2}$. In fact, $(\frac{5}{6}, \frac{1}{2})$ is a Nash equilibrium: from the above argument, Bob's payout is maximal for the given value of p. Similarly, Alice's payout is maximal for the given value of q. So, neither player has an incentive to change their strategies.

Nonexample. Nash equilibrium may not exist if the set of choice is infinite and non-compact (for comparison, the k-simplex is compact in $\mathbb{R}^{k+1}$). Consider the 2-player game in which the two players each pick a real positive number, and the player who picked the bigger number wins. There is no equilibrium since any player can always pick a number bigger than the previously played numbers. Another example is if the players must pick a real number in the interval $[0, 5)$, and the player with the bigger number wins. For this game, Nash equilibrium does not exist, and this is fine since the set of choices is non-compact.

Proof of Existence of Nash equilibrium. The proof of the existence of Nash equilibrium is clever, but not hard to follow. The only tricky part is actually coming up with a continuous function from a space of strategy combinations to itself such that a fixed point of the function is a Nash equilibrium. First, we define a function that measures improvement by switching to a pure strategy:

$$\phi_{i\alpha}(\mathbb{S}) = \max(0, p_i(s_1, \ldots, \pi_{i\alpha}, \ldots, s_n) - p_i(\mathbb{S})).$$

Notice that the improvement function is positive if player i increases his or her expected payout by switching to the α^{th} pure strategy, and zero otherwise.

We claim that $\phi_{i\alpha}(\mathbb{S}) = 0$ for all α if and only if player i has no better strategy i.e., there is no strategy with a greater expected payout if the other players keep their strategies fixed. Suppose that player i has no better strategy. By definition, any strategy available to player i must have a payout less or equal to $p_i(\mathbb{S})$. In particular, any pure strategy must be equally good or worse than the current strategy for player i. Hence, the improvement function is exactly 0. Conversely, suppose that $\phi_{i\alpha}(\mathbb{S}) = 0$ for all α. Our current strategy s_i is at least as profitable as any pure strategy $\pi_{i\alpha}$. Linearity of p_i tells us that s_i is as profitable as any mixed strategy. Hence, player i cannot increase p_i by changing his strategy, so there is no better strategy for player i.

It follows from the definition of Nash equilibrium that $\phi_{i\alpha}(\mathbb{S}) = 0$ for all i, α if and only if $\mathbb{S}$ is a Nash equilibrium.

With the improvement functions $\phi_{i\alpha}$, we define the continuous map for Brouwer's fixed point theorem. For each component s_i of $\mathbb{S}$, define a modification s_i' by

$$s_i' = \frac{s_i + \sum_\alpha \phi_{i\alpha}(\mathbb{S})\pi_{i\alpha}}{1 + \sum_\alpha \phi_{i\alpha}(\mathbb{S})}.$$

This gives a map $T : \mathbb{S} \mapsto \mathbb{S}' = (s_1', s_2', \ldots, s_n')$. The set of strategy combinations is homeomorphic to $\overline{B^m}$ for some $m > 0$, and T is a continuous map. Hence, Brouwer's fixed point theorem gives us a fixed point of T. We claim that a strategy combination is a fixed point of T if and only if it is a Nash equilibrium. To show this, we will prove that a strategy combination is a fixed point of T if and only if $\phi_{i\alpha}(\mathbb{S}) = 0$ for all i, α.

The right-to-left direction is easy: setting $\phi_{i\alpha}(\mathbb{S}) = 0$ for all i, α means that $s_i' = s_i$ so $\mathbb{S} = T(\mathbb{S})$ and $\mathbb{S}$ is a fixed point of T.

Now we prove the other direction. For any player i, we can order the pure strategies currently in use (pure strategies $\pi_{i\alpha}$ such that $c_{i\alpha'} > 0$) in order of increasing payout. Hence, we can find a "least profitable" strategy $\pi_{i\alpha'}$ such that $c_{i\alpha} > 0$. The linearity of p_i implies that

$$p_i(s_1, \ldots, \pi_{i\alpha'}, \ldots, s_n) \leq p_i(\mathbb{S}),$$

or equivalently, $\phi_{i\alpha'}(\mathbb{S}) = 0$. Our goal is to show that all pure strategies that player i uses are equally profitable as $\pi_{i\alpha'}$.

Let $\mathbb{S} = (s_1, \ldots, s_n)$ be a fixed point of T. Then,

$$s_i = \frac{s_i + \sum_\alpha \phi_{i\alpha}(\mathbb{S})\pi_{i\alpha}}{1 + \sum_\alpha \phi_{i\alpha}(\mathbb{S})}, \tag{6.1}$$

where we have just replaced s_i' with s_i. Because $\phi_{i\alpha'}(\mathbb{S}) = 0$ and the $\pi_{i\alpha}$ are linearly independent vectors, for the proportion of $\pi_{i\alpha'}$ in s_i to remain the same under the map T, the denominator of the fraction in (6.1) must be 1. This implies that $\phi_{i\alpha}(\mathbb{S}) = 0$ for all α, and by extension, for all i, α. Thus, we have proved that the fixed point of the map $T : \mathbb{S} \mapsto \mathbb{S}'$ is a Nash equilibrium. This proves the existence of a Nash equilibrium for n-player games in which each player has finitely many pure strategies. ■