Kakutani's Fixed-Point Theorem

And its Applications to Game Theory

Definition. *A **convex combination** of the set $V = \{\boldsymbol{v}_1, \boldsymbol{v}_2, \cdots, \boldsymbol{v}_m\}$ is a linear combination such that the coefficients are all non-negative and sum to 1. The set of all convex combinations of V*

$$C = \{\sum_{i=1}^{m} \lambda_i \boldsymbol{v}_i \mid \lambda_i \geq 0 \text{ for all } i \text{ and } \sum_{i}^{m} \lambda_i = 1\}$$

*is called the **convex hull** of V.*

Definition. *A set $C \subseteq \mathbb{R}^n$ is said to be **convex** when it contains all of its convex combinations. Equivalently, given any 2 points $c_1, c_2 \in C$, it follows that $\lambda c_1 + (1-\lambda)c_2 \in C$ for all $\lambda \in [0,1]$.*

The points $\lambda c_1 + (1-\lambda)c_2 \in C$ are just the points on the line segment between c_1 and c_2. In other words, the 1-simplex of these two. For example, the point when $\lambda = 0$ is c_2 and when $\lambda = 1$ we have c_1. When $\lambda = \frac{1}{2}$ we have the point exactly in between c_1 and c_2, and when $\lambda = \frac{1}{3}$ we have the point 1 third the way from c_2 to c_1 (closer to c_2) and so on.

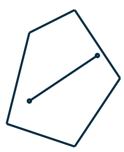

Figure 1: A convex subset of $\mathbb{R}^2$(left) and a non-convex one (right).

Definition. *A the set $V = \{\boldsymbol{v}_1, \boldsymbol{v}_2, \cdots, \boldsymbol{v}_m\}$ is said to be **affine independent** when $\sum_{i=1}^{n} \lambda_i \boldsymbol{v}_i = 0$ and $\sum_{i=1}^{n} \lambda_i = 0$ implies $\lambda_i = 0$ for all i. Equivalently, when the set $\{\boldsymbol{v}_2 - \boldsymbol{v}_1, \boldsymbol{v}_2 - \boldsymbol{v}_1, \cdots \boldsymbol{v}_m - \boldsymbol{v}_1\}$ is linearly independent.*

Definition. *An n-**simplex** is the convex hull of an affine independent set of $n+1$ vectors. The **standard** n-**simplex** is the convex hull of the standard basis in $\mathbb{R}^{n+1}$*

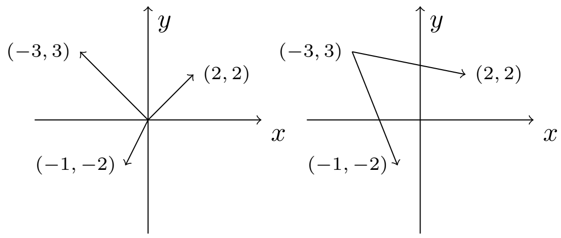

Figure 2: The vectors $(-3,3)$, $(-1,-2)$ and $(2,2)$ being affine independent can be thought of as meaning that the vectors are linearly independent from the point of view of $(-3,3)$.

Definition. *If S is an n-simplex whose vertices are the set $V = \{\boldsymbol{v}_1, \boldsymbol{v}_2, \cdots, \boldsymbol{v}_{n+1}\}$, then a **face** of S is a simplex whose vertices are a subset of V. A face is a **proper face** when its vertices are a strict subset of V.*

Definition. *Let S be an n-simplex. A **triangulation** of S is a finite collection $T = \{S_1, S_2, S_3, \cdots S_m\}$ such that*

1. *$S_i \subseteq S$ for all $i \in \{1, 2, \cdots, m\}$*
2. *Each S_i is an n-simplex*

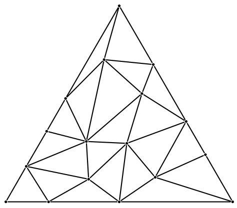

Figure 3: A triangulation of a standard 2-simplex

3. *If $i \neq j$, then $S_i \cap S_j$ is empty, or is an m-simplex which is a proper face shared by S_i and S_j*
4. $\cup_{i=1}^{m} S_i = S$

Intuitively, a triangulation of an n-simplex is just cutting up the simplex into smaller simplices. Here is a useful theorem about n-simplices being closed sets:

Theorem. *The convex hull of a set of vectors $V = \{\boldsymbol{v}_1, \boldsymbol{v}_2, \cdots, \boldsymbol{v}_m\}$ is a closed set.*

Proof. Let $(\mathbf{x}_n) \to \mathbf{x}_0$ be a sequence in the convex hull of V. Then for each term we have $\mathbf{x}_n = \sum_{i=1}^{m} \lambda_i^n \mathbf{v}_i$ where each $\lambda_i^n \geq 0$ and $\sum_{i=1}^{m} \lambda_i^n = 1$ for every $n \in \mathbb{N}$. Then,

$$\mathbf{x}_0 = \lim_{n\to\infty} \mathbf{x}_n = \lim_{n\to\infty} \sum_{i=1}^{m} \lambda_i^n \mathbf{v}_i = \sum_{i=1}^{m} \lim_{n\to\infty} (\lambda_i^n) \mathbf{v}_i$$

Since each λ_i^n is in the closed set $[0,1]$, we know there exists a convergent subsequence for each i. We call the limit of this sequence λ_i^0. Note that λ_i^0 is non-negative. The corresponding subsequence of $(\mathbf{x}_n)$ will still converge to $\mathbf{x}_0$ since it is already convergent. We assume for simplicity that (λ_i^n) itself converges. We've shown that $\mathbf{x}_0$ is a linear combination of V with coefficients λ_i^0. All that is left to do is to show that they sum to 1. Indeed,

$$\sum_{i=1}^{m} \lambda_i^0 = \sum_{i=1}^{m} \lim_{n\to\infty} \lambda_i^n = \lim_{n\to\infty} \sum_{i=1}^{m} \lambda_i^n = \lim_{n\to\infty} 1 = 1$$

Thus, $\mathbf{x}_0$ is a convex combination of V, and the convex hull is closed. □

Brouwer's Fixed Point Theorem

In this section, we prove Brouwer's fixed point theorem for the case of continuous functions from an n-simplex to itself. This is a combinatorial proof, and so we will introduce some background material first.

Definition. *A graph G is a pair (V, E) where V is a (finite) set of vertices, and E a set of pairs of vertices from V. The degree of a vertex is the number of pairs in E it appears in.*

Intuitively, a graph is a set of vertices, which we will from now on call nodes so as to not confuse with the vertices of simplices, and edges connecting the nodes in different ways. The degree of a node is the amount of nodes it is connected to.

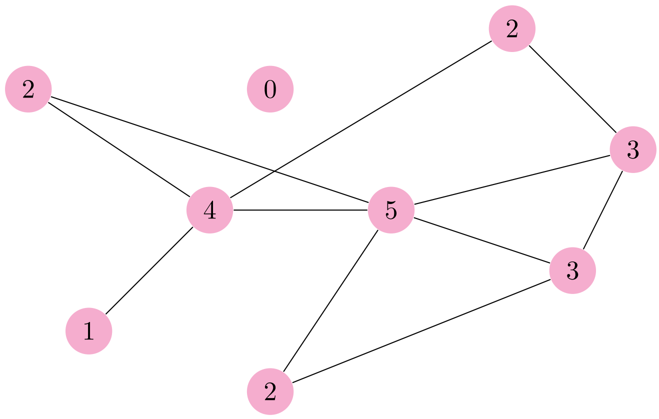

Figure 4: A graph with each node labelled with its degree

There are many applications of these graphs. For example, they may represent social connections on social media such as facebook, where each account is a node, and nodes are connected when the accounts they represent are friends. There are many other things that can be done, for example, we can also associate 'weights' with each edge to give a sense of how strongly two nodes are connected, or we can also have directed graphs where the edges have a direction to them. This can be useful to represent connections which are not necessarily mutual. For example, on twitter, you can follow people who may not follow you or vice versa, which is different from social sites like facebook where a friend connection is mutual.

Lemma (Handshaking Lemma)**.** *For any graph G, the number of nodes of odd degree is even.*

Proof. Summing up the degrees of all nodes counts over each edge twice. This is because each edge contributes to the degree of the two nodes it connects. This means, the sum of all the degrees is twice the number of edges. So it is even. This can only be true if there are evenly many nodes of odd degree. □

Informally, if you have some number of people at a party shaking each other's hands (not counting one person shaking the same hand twice), there is an even number of people who shook the hands of an odd number of other people.

Definition (Sperner Colouring)**.** *Let S be an n-simplex and T a triangulation of S. Let $V = \{v_1, v_2, \cdots, v_{n+1}\}$ be the set of vertices of simplices in T. Then a function $f : V \to \{1, 2, \cdots, n, n+1\}$ is a Sperner colouring when the following hold:*

1. *For each $v_k \in V$, $f(v_k) = k$.*
2. *If u is a vertex of a simplex in T and is located on a subface of S formed by the vertices $v_{n_1}, v_{n_2}, \cdots, v_{n_k}$, then $f(u) \in \{n_1, n_2, \cdots, n_k\}$*

The numbers in the set $\{1, 2, \cdots, n, n+1\}$ can instead be thought of as a set of $n+1$ colours, with which we are colouring the vertices (hence the name). See figure 5.

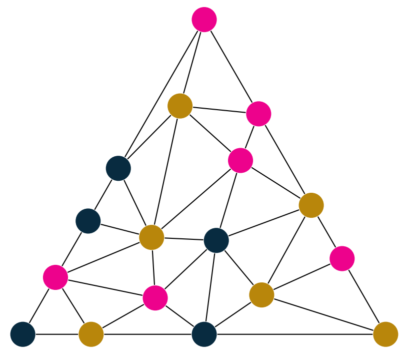

Figure 5: A Sperner colouring of a 2-simplex. The vertices of the triangulation on each edge of the large triangle are only coloured with the colours of the vertices of that edge. Moreover, each vertex of the large triangle is of a different colour.

Lemma (Sperner's Lemma)**.** *Let T be a triangulation of an n-simplex S, where V denotes the set of vertices of simplices in the triangulation. Let $f : V \to \{1, 2, \cdots, n+1\}$ be a Sperner colouring. Then there exists at least one sub-simplex $S' \in T$ which is coloured with all $n+1$ colours. Moreover, there exist an odd number of such simplices.*

Proof. The case for a 1-simplex is easy. Let T be a triangulation of this simplex. By the Sperner colouring rules, one end of this simplex is of colour 1, and the other of colour 2. Moving from one end sequentially through the vertices of the triangulation, the colours of the vertices must switch between 1 and 2 an odd number of times in order to start at 1 and end at 2.

For the induction step, assume that the lemma holds on any $n-1$-simplex, and consider an n-simplex S with vertices $s_1, s_2, \cdots, s_{n+1}$. Let T be a triangulation of the n-simplex, and let each vertex in T be coloured with the colours $\{1, 2, \cdots, n+1\}$ in accordance with the rules of a Sperner colouring and for simplicity, assume each s_k is coloured k for each $1 \leq k \leq n+1$. We construct a graph G as follows. Each simplex in T is represented by a node, and there is one more node representing the (closed) region outside the simplex. Clearly, given any two distinct nodes in G, the intersection of the regions represented by them is either empty, or an m-simplex where $m < n$. Two nodes will be connected precisely when this intersecion is a simplex whose vertices are coloured with all the colours $\{1, 2, \cdots, n\}$ (but not $n+1$).

Note that only simplices in T which have a face on the boundary of S can be can share a region with the exterior. More importantly, only the $(n-1)$-simplices on the face of S with vertices $s_1, s_2, \cdots, s_n$ can be coloured with the colours required for a connection. This is a consequence of the rules of a Sperner colouring. Then by the induction hypothesis, there are an odd number of nodes connected to the external node. It follows from the handshaking lemma that there exist an odd number of nodes representing simplices in T which are connected to oddly many other nodes.

Now consider a simplex in T represented by a node of degree 2 or greater. This simplex has $n+1$ vertices, and n of them are coloured with the colours $\{1, 2, \cdots, n\}$. The remaining vertex cannot be coloured $n+1$, otherwise its node would not have a degree of at least 2. So it is coloured one of the colours $\{1, 2, \cdots, n\}$. Then by simple counting, exactly two of the faces of this simplex are coloured with each of $\{1, 2, \cdots, n\}$, and so its node is connected to exactly 2 others. So every simplex in T has a node of degree either 0, 1, or 2. Since there must be an odd number of simplices in T with nodes of odd degree, there is a simplex with node of degree 1. It is easy to see that this simplex is fully coloured with the colours $\{1, 2, \cdots, n+1\}$ which is what we wanted to show. □

This lemma will be quite useful in proving Brouwer's fixed point theorem. The general idea is to take a sequence of triangulations where the simplices become arbitrarily small on a standard n-simplex, and define a Sperner colouring on the vertices of each triangulation. Then we have a sequence of simplices in each triangulation which are fully coloured with all the colours. The vertices of these simplices form sequencs which approach a single point due to the fact that the vertices of a simplex in the triangulations become arbitralily close to one another. This point turns out to be the required fixed point.

Theorem (Brouwer's Fixed Point Theorem for n-simplices)**.** *Let Δ_n be the standard n-simplex in $\mathbb{R}^{n+1}$, and let $f : \Delta_n \to \Delta_n$ be a continuous function. There exists a point $x \in \Delta_n$ such that $x = f(x)$.*

Proof. Let (T_i) be a sequence of triangulations of Δ_n where the diameter's approach 0, and let $V_i = \{\mathbf{v}_1^i, \mathbf{v}_2^i, \mathbf{v}_3^i, \cdots \mathbf{v}_{n_i}^i\}$ be the set of vertices for all simplices in the triangulation. For this proof, we will denote that kth entry of a vector $\mathbf{v}_j^i$ as $[\mathbf{v}_j^i]_k$.

We will now define for each T_i a function $\lambda_i : V_i \to \{1, 2, \cdots, n+1\}$ as follows. Consider an arbitrary vertex $\mathbf{v} \in V_i$. By the definition of the standard n-simplex, the entries of $\mathbf{v}$ sum up to 1. Similarly for the vector $f(\mathbf{v})$. Clearly then, there is an entry of $\mathbf{v}$ which is greater than or equal to the corresponding entry in $f(\mathbf{v})$ and strictly positive. The index of this entry will be the image of $\mathbf{v}$ under the function λ_i. It is easy to verify that this is a Sperner colouring. Each vertex of the large simplex Δ_n only has one non-zero entry, and the index of this entry will be its image under λ_i. Clearly, each of these vertices will have a different colour. For simplicity, let these vertices be $s_1, s_2, \cdots, s_{n+1}$ where each s_k is the kth standard basis vector of $\mathbb{R}^{n+1}$, and so $\lambda_i(s_k) = k$ for each. The second rule is fulfilled by the observation that if a vertex of the triangulation is on a subface of Δ_n with vertices $s_{n_1}, s_{n_2}, \cdots, s_{n_k}$, then it is a convex combination of these and so only the entries indexed $n_1, n_2, \cdots, n_k$ can be non-zero. So the colour of this vertex is one of $n_1, n_2, \cdots, n_k$ as desired.

Then applying Sperner's lemma, there exists an n-simplex $S_i \in T_i$ for each triangulation T_i such that the vertices of S_i are $\{s_1^i, s_2^i, \cdots s_n^i, s_{n+1}^i\} \subseteq V_i$ and $\lambda_i(s_k^i) = k$ for all $k \in \{1, 2, 3, \cdots, n+1\}$.

Consider the kth vertex of each S_i, and we have the sequence $(s_k^i)_{i\in\mathbb{N}}$ for each $k \in \{1, 2, 3, \cdots, n+1\}$. That is, we have a sequence of vertices for each colour. Each is bounded within the compact set Δ_n, so by the Bolzano-Weierstrass theorem, contains a convergent subsequence. For simplicity, assume each $(s_k^i)_{i\in\mathbb{N}}$ itself is the convergent subsequence, and denote the limit s_k^0. However, note that the triangulations become arbitrarily fine, so the sequences all approach the same limit. That is, $s_1^0 = s_2^0 = s_3^0 = \cdots = s_{n+1}^0$. We call this point $\mathbf{x}^0$, and we will now show that this is the desired fixed point of the continuous function f.

By continuity, because $[s_k^i]_k \geq [f(s_k^i)]_k$ for every triangulation T_i, it follows that the kth entry of the limit is also at least as large as the kth entry of its image under f. So $[\mathbf{x}^0]_k \geq [f(\mathbf{x}^0)]_k$. This is true for every entry of $\mathbf{x}^0$ since it is the limit of $(s_k^i)_{i\in\mathbb{N}}$ for every k. But to ensure that the sum of the entries of both $\mathbf{x}$ and $f(\mathbf{x}^0)$ sum to 1, the inequalities must in fact be equalities. That is, $[\mathbf{x}^0]_k = [f(\mathbf{x}^0)]_k$ for each $k \in \{1, 2, \cdots, n+1\}$. Thus, $\mathbf{x}^0 = f(\mathbf{x}^0)$ as desired. □

This is a special case of the foull Brouwer fixed point theorem. The general statement is as follows:

Theorem (Brouwer's Fixed Point Theorem). *If K is a non-empty, compact, and convex subset of $\mathbb{R}^n$, and $f : K \to K$ is continuous, then there exists $x \in K$ such that $x = f(x)$.*

Kakutani's Theorem

Theorem (Kakutani's Theorem for n-simplices)**.** *Let $\Phi(x) : \Delta_n \to \mathcal{P}(\Delta_n)$ be an upper semi-continuous function such that $\Phi(x)$ is a non-empty, convex set for all $x \in \Delta_n$. Then there exists $x_0 \in \Delta_n$ such that $x_0 \in \Phi(x_0)$*

Proof sketch.

1. Construct a series of triangulations (T_i) where the mesh goes to 0
2. Define a function φ_i on each triangulation which maps each vertex of T_i to some element of its image under Φ
3. Extend the function linearly over each sub-simplex. Appeal to Brouwer's fixed point theorem and denote the fixed point as $\mathbf{x}^i$.
4. The sequence of $\mathbf{x}^i$ is in a compact set, so denote the limit $\mathbf{x}^0$.
5. Each $\mathbf{x}^i$ is a convex combination of the vertices of the simplex in T_i it is contained in. For each T_i, number the vertices of this simplex with $1, 2, \cdots, n+1$. The vertices numbered 1 form a sequence, and so do those numbered 2, and so on.
6. We also define for each sequence of vertices a sequence of the image of each vertex under φ_i. We have $n+1$ sequences of vertices, and their corresponding images under φ_i. Each has convergent subsequence.
7. As $i \to \infty$, the mesh approaches 0 so each sequence of vertices in fact approaches $\mathbf{x}^0$.
8. The sequence of vertices approaches the limit $\mathbf{x}^0$, and the sequence of the images of the vertices under φ_i approaches a limit. Each term of the latter is contained within the image of each term of the former under Φ. By upper semi-continuity, the limit of the images of the vertices under φ_i is contained in the image of $\mathbf{x}^0$ under Φ.
9. Show that $\mathbf{x}^0$ is in fact a convex combination of the limits of the images of the vertices under φ_i.
10. The result follows from the fact that Φ maps to convex sets

Proof. We again have a sequence of triangulations $(T_i)_{i\in\mathbb{N}}$ on the n-simplex Δ_n such that the mesh approaches 0. Let $V_i = \{\mathbf{v}_1^i, \mathbf{v}_2^i, \mathbf{v}_3^i, \cdots \mathbf{v}_{n_i}^i\}$ be the vertices in V_i. For each triangulation T_i, we will define a function $\varphi_i : \Delta_n \to \Delta_n$ as follows.

For each vertex $\mathbf{v}_j^i \in V_i$, we will define $\varphi_i(\mathbf{v}_j^i)$ as any of the vectors in $\Phi(\mathbf{v}_j^i)$. That is, $\varphi_i(\mathbf{v}_j^i) \in \Phi(\mathbf{v}_j^i)$ for all $\mathbf{v}_j^i \in V_i$. Next we extend φ_i linearly over each simplex in T_i. That is, if $\mathbf{x}$ is in the simplex with vertices $\mathbf{u}_1^i, \mathbf{u}_2^i, \cdots, \mathbf{u}_{n+1}^i$, then

$$\mathbf{x} = \sum_{j=1}^{n} \alpha_j^i \mathbf{u}_j^i$$

Where $\alpha_k^i \geq 0$ for all $k \in \{1, 2, 3, \cdots n+1\}$ and $\sum_{j=1}^{n} \alpha_j^i = 1$. Then $\varphi_i(\mathbf{x})$ is defined as

$$\varphi_i(\mathbf{x}) = \varphi_i\left(\sum_{j=1}^{n+1} \alpha_j^i \mathbf{u}_j^i\right) = \sum_{j=1}^{n+1} \alpha_j^i \varphi_i(\mathbf{u}_j^i)$$

See that each $\varphi_i : \Delta_n \to \Delta_n$ is a continuous function between compact sets. So by Brouwer's fixed point theorem, there is a fixed point for each φ_i which we will denote $\mathbf{x}^i$. Such a point is defined for each T_i, so in fact we have a sequence of points $(\mathbf{x}^i)_{i\in\mathbb{N}}$ which is clearly bounded in the compact set Δ_n. So there is a convergent subsequence. To avoid having to use more complicated indexing, we assume $(\mathbf{x}^i)_{i\in\mathbb{N}}$ is indeed such a subsequence and $\lim\limits_{i\to\infty} \mathbf{x}^i = \mathbf{x}^0$. We claim this is the desired fixed point.

Each $\mathbf{x}^i$ is contained in a simplex of T_i with vertices $\mathbf{w}_1^i, \mathbf{w}_2^i, \cdots, \mathbf{w}_{n+1}^i$. So we also have bounded sequences $(\mathbf{w}_k^i)_{i\in\mathbb{N}}$ for each $k \in \{1, 2, \cdots, n+1\}$, and these also contain convergent subsequences. We again assume they are themselves are convergent for simplicity and for each k, $\lim\limits_{i\to\infty} \mathbf{w}_k^i = \mathbf{w}_k^0$. Moreover, $\mathbf{x}^i$ may be expressed as a convex combination

$$\mathbf{x}^i = \sum_{j=1}^{n+1} \lambda_j^i \mathbf{w}_j^i$$

With the usual restriction on each λ_j^i. The image of each $\mathbf{w}_j^i$ under φ_i will be denoted $\mathbf{y}_j^i$. Because the sequence $(\mathbf{w}_k^i)_{i\in\mathbb{N}}$ converges, the sequence $(\mathbf{y}_k^i)_{i\in\mathbb{N}}$ will also converge since φ_i is continuous. Let this limit be $\mathbf{y}_k^0$. Similarly, let $(\lambda_k^i)_{i\in\mathbb{N}} \to \lambda_k^0$.

To summarise, we have defined a function φ_i for each tringulation T_i on Δ_n and defined its fixed point as $\mathbf{x}^i$. This fixed point must be contaied in some simplex of the triangulation T_i, and we numbered the vertices from 1 to $n+1$. The vertices numbered k form a sequence which we denote $(\mathbf{w}_k^i)_{i\in\mathbb{N}}$. Similarly, we have defined for the sequence of vertices numbered k a sequence of their images under its corresponding function φ_i, and we denote this $(\mathbf{y}_k^i)_{i\in\mathbb{N}}$. Since each fixed point is a convex combination of the vertices of the simplex in T_i it is located in, we have defined sequences of the coefficiens as well. Putting all these limits together gives the following

$$\begin{aligned}
\mathbf{x}^0 &= \lim_{i\to\infty} \mathbf{x}^i = \lim_{i\to\infty} \varphi_i(\mathbf{x}^i) = \lim_{i\to\infty} \varphi_i\left(\sum_{j=1}^{n+1} \lambda_j^i \mathbf{w}_j^i\right) = \lim_{i\to\infty}\left(\sum_{j=1}^{n+1} \lambda_j^i \varphi_i(\mathbf{w}_j^i)\right) \\
&= \lim_{i\to\infty}\left(\sum_{j=1}^{n+1} \lambda_j^i \mathbf{y}_j^i\right) = \sum_{j=1}^{n+1} \lim_{i\to\infty} \lambda_j^i \lim_{i\to\infty} \mathbf{y}_j^i = \sum_{j=1}^{n+1} \lambda_j^0 \mathbf{y}_j^0
\end{aligned}$$

Moreover,

$$\sum_{j=1}^{n+1} \lambda_j^0 = \sum_{j=1}^{n+1} \lim_{i\to\infty} \lambda_j^i = \lim_{i\to\infty} \sum_{j=1}^{n+1} \lambda_j^i = \lim_{i\to\infty} 1 = 1$$

Each λ_k^0 is in the set $[0,1]$ since they are limits of sequences in this closed set. So $\mathbf{x}^0$ is a convex combination of the set $\{\mathbf{y}_1^0, \mathbf{y}_2^0, \cdots, \mathbf{y}_{n+1}^0\}$.
Going back to the sequence $(\mathbf{w}_k^i)_{i\in\mathbb{N}}$, we said that this conveges to $\mathbf{w}_k^0$ for each k. We defined these to be the vertices of the simplex in V_i containing the fixed point $\mathbf{x}^i$. The simpleces in the triangulations all approach 0, so it is not hard to show that each of these sequence in fact converges to the limit of $(\mathbf{x}^i)_{i\in\mathbb{N}}$. That is, $\mathbf{x}^0$. So we have for each triangulation k the following

$$(\mathbf{w}_k^i)_{i\in\mathbb{N}} \to \mathbf{x}^0$$

$$(\mathbf{y}_k^i)_{i\in\mathbb{N}} \to \mathbf{y}_k^0$$

$$\mathbf{y}_k^i = \varphi_i(\mathbf{w}_k^i) \in \Phi(\mathbf{w}_k^i)$$

From the upper semi-continuity of Φ, it follows that $\mathbf{y}_k^0 \in \Phi(\mathbf{x}^0)$ for each k. But one of the premises is that Φ maps to convex sets. So any convex combination of $\{\mathbf{y}_1^0, \mathbf{y}_2^0, \cdots, \mathbf{y}_{n+1}^0\}$ must also be in $\Phi(\mathbf{x}^0)$. We have already shown that $\mathbf{x}^0$ is such a convex combination, so we conclude that $\mathbf{x}^0 \in \Phi(\mathbf{x}^0)$ as desired. $\square$

Like Brouwer, this is also just a special case of the full Kakutani fixed point theorem. The full theorem applies to arbitrary non-empty, compact and convex subsets of Euclidean space and not just simplices.

Theorem (Kakutanis fixed point theorem)**.** *Let S be a non-empty, compact and convex subset of Euclidean space, and let $\Phi : S \to \mathcal{P}(S)$ be a set valued function such that*

1. *Φ is upper-semicontinuous*

2. *$\Phi(s)$ is non-empty and convex for all $s \in S$*

Then there exists $s_0 \in S$ such that $s_0 \in \Phi(s_0)$.

This more general statement is the one we will need in order to prove the existence of Nash equilibria for mixed extensions.

Review of Game Theory Terms

Definition. *A finite, n-person normal-form game is a tuple (N, A, u) where*

1. *N is a set of n players indexed by i*
2. *$A = \prod_{i=1}^{n} A_i$ is the set of action profiles, where each A_i is the set of actions available to player i.*
3. *$u = (u_1, u_2, \cdots u_n)$ is a utility function $u : A \to \mathbb{R}^n$, where each $u_i : A \to \mathbb{R}$ is a real valued utility function for player i*

Definition. *The mixed extension of the n-person normal-form game (N, A, u) is the tuple (N, S, U) where*

1. *$S = \prod_{i=1}^{n} S_i$ where S_i is the set of all probability distributions over A_i*
2. *$U = (U_1, U_2, \cdots, U_n)$ where each U_i is defined as a function $U_i : S \to \mathbb{R}$*

$$U_i(s) = \sum_{\boldsymbol{a} \in A} u_i(\boldsymbol{a}) \prod_{i=1}^{n} s_i(\boldsymbol{a}_i)$$

Definition. *The best response function for player i is a set-valued function $B_i : S_{-i} \to \mathcal{P}(S_i)$ where*

$$B_i(s_{-i}) = \{s_i^* \in S_i \mid U_i(s_i^*, s_{-i}) \geq U_i(s_i, s_{-i}) \text{ for all } s_i \in S_i\}$$

A minor detail is that the best response does indeed exist for every $s \in S$. Since U_i is defined on a compact set, and is continuous, the extreme value theorem applies and the best response does indeed exist for all points in the domain.

We extend the best reponse function and define $B : S \to \mathcal{P}(S)$

$$B(s) = \prod_{i=1}^{n} B_i(s_{-i})$$

Where $\prod$ denotes the cartesian product. The function B takes in mixed strategies, and returns a the optimal strategies that each player should have played for the highest expected utility. This motivates the following (re)definition:

Definition. *A mixed strategy $s \in S$ is said to be a **Nash equilibrium** when $s \in B(s)$.*

This is because $B(s)$ is the cartesian product of all the players' individual best response functions. So $s \in B(s)$ means that for each player, $s_i \in B_i(s_{-i})$ and so the strategy profile for each player is already the optimal strategy given the other player's strategies, which is how Nash equilibria were defines before.

It is now more clear how Kakutani's fixed-point theorem applies to game theory. The cartesian product of all the players' best response functions is a set valued function taking in mixed strategies and outputting a set of of lists whose entries are the possible optimal strategies each player could have played. If each player has already played an optimal strategy such that they cannot expect a strictly higher utility by changing, then the input of the function B will be in its output. So the Nash Equilibrium is a fixed point of a set-valued function. So if we can just verify the conditions under which Kakutani's fixed point theorem applies, we can easily prove the existence of Nash equilibria of mixed extensions of finite n-person normal form games.

It has bee argued that the set of mised strategies for player i can be seen as a standard simplex. Say the player has $n + 1$ possible actions. Then each action can be thought of as one of the standard basis vectors in $\mathbb{R}^{n+1}$. Then extend this idea to mixed strategy profiles by weighing each vector by it probability so that we have a set of points in $\mathbb{R}^{n+1}$ which is a linear combination of the basis vectors such that the coefficients of the combinations are non-negative and sum to 1; this is by definition a standard n-simplex.

All that is left to show is that :

1. S, the cartesian product of each S_i is compact and convex,
2. B is upper semi-continuous,
3. $B(s)$ is convex for all $s \in S$.

And then we will know that Kakutani's theorem applies.

Lemma. *The cartesian product of finitely many convex sets is also convex.*

Proof. Let $(x_1, x_2, \cdots, x_n), (x'_1, x'_2, \cdots, x'_n) \in X_1 \times X_2 \times \cdots \times X_n$ where each X_i is a convex subset of some euclidean space of arbitrary dimension. Consider the point $\lambda(x_1, \cdots, x_n) + (1-\lambda)(x'_1, \cdots, x'_n)$ where $\lambda \in [0,1]$. Then see that

$$\lambda(x_1, x_2, \cdots, x_n) + (1-\lambda)(x'_1, x'_2 \cdots, x'_n) = (\lambda x_1 + (1-\lambda)x'_1, \cdots, \lambda x_n + (1-\lambda)x'_n)$$

Since each X_i is convex, $\lambda x_i + (1-\lambda)x'_i \in X_i$ for each i, and so

$$(\lambda x_1 + (1-\lambda)x'_1, \lambda x_2 + (1-\lambda)x'_2, \cdots, \lambda x_n + (1-\lambda)x'_n) \in X_1 \times X_2 \times \cdots \times X_n$$

Thus, $X_1 \times X_2 \times \cdots \times X_n$ is convex. □

Lemma. *The cartesian product of finitely many closed sets is also closed.*

Proof. Let $X_1, X_2, \cdots, X_k$ be closed sets and let $(x_1^n, x_2^n, \cdots, x_k^n) \to (x_1, x_2, \cdots, x_k)$ be a sequence where $(x_1^n, x_2^n, \cdots, x_k^n) \in X_1 \times X_2, \times \cdots \times X_k$ for all $n \in \mathbb{N}$. Then we have for each $1 \leq i \leq k$, we have the sequence $(x_i^n) \to x_i$ where $x_i^n \in X_i$ for all n. Then since X_i is closed, it must be be true that $x_i \in X_i$. This is true for each i, so it follows that $(x_1, x_2, \cdots, x_k) \in X_1 \times X_2, \times \cdots \times X_k$. Thus, the cartesian product is indeed closed. □

This completes the first requirement which says that we want S to be compact. By the Heine Borel theorem this is the same as a set being closed and bounded. It is clear that the standard n-simplices are bounded sets in Euclidian space, and so their cartesian product should be bounded in its Euclidean space. A more general result about the cartesian product of compact spaces being compact is Tychonoff's theorem which holds for arbitrary topological spaces, although this result looks far too technical.

We now go on to show that B is an upper-semi continuous function.

Lemma. *The best response function $B(s) : S \to S$ is upper semi-continuous.*

The proof for this is quite heavy in notation and confusing. The following is a proof sketch:

1. Consider sequences in S, $(r^n) \to r^0$ and $(s^n) \to s^0$ be sequences in S such that $r^n \in B(s^n)$ for all $n \in \mathbb{N}$. Assume for contradtiction that $r^0 \notin B(s^0)$
2. Using the fact that U_i is continuous, show that somewhere along the sequence r^n stopped being the best response to s^n which would be a contradiction.

Proof. Let $(r^n) \to r^0$ and $(s^n) \to s^0$ be sequences in S. Let it also be true that $r^n \in B(s^n)$ for all $n \in \mathbb{N}$. Assume for contradiction that B is not upper semi-continuous, so $r^0 \notin B(s^0)$. Then for some i, we have $r_i^0 \notin B_i(s_{-i}^0)$. Then let $r'_i \in B_i(s_{-i}^0)$, so there exists $\varepsilon > 0$ such that

$$U_i(r'_i, s_{-i}^0) > U_i(r^0, s_{-i}^0) + 3\varepsilon$$

Morover, we can make $U_i(r'_i, s_{-i}^n)$ arbitrariliy close to $U_i(r'_i, s_{-i}^0)$ by bringing s_{-i}^n sufficiently close to s_{-i}^0. So for sufficiently large n,

$$U_i(r'_i, s_{-i}^n) > U_i(r'_i, s_{-i}^0) - \varepsilon$$

Similarly, $U_i(r^n, s_{-i}^n)$ can be made arbitrarily close to $U_i(r^0, s_{-i}^0)$ with sufficiently large n. So

$$U_i(r^0, s_{-i}^0) > U_i(r^n, s_{-i}^n) - \varepsilon$$

So putting these 3 inequalities together gives

$$U_i(r'_i, s_{-i}^n) > U_i(r'_i, s_{-i}^0) - \varepsilon > U_i(r^0, s_{-i}^0) + 2\varepsilon > U_i(r^n, s_{-i}^n) + \varepsilon$$

But we had a premise that $r^n \in B(s^n)$, and thus $r_i^n \in B_i(s_{-i}^n)$ for all n. So this is a contradiction and so we conclude that B is upper semi-continuous as desired. □

This completes the second requirement. The third is easier to show and is quite straightforward.

Lemma. *For all $s \in S$, the set $B(s)$ is convex*

Proof. Let $s_k \in S_k$ be such that $s_k = \lambda s_{k_1} + (1-\lambda)s_{k_2}$ where $s_{k_1}, s_{k_2} \in S_k$ and $0 \leq \lambda \leq 1$. We will first demonstrate that $U_i(s_k, s_{-k}) = \lambda U_i(s_{k_1}, s_{-k}) + (1-\lambda)U_i(s_{k_1}, s_{-k})$ This is easy to show

$$\begin{aligned}
U_i(s_1, \cdots, s_k, \cdots s_n) &= U_i(s_1, \cdots, \lambda s_{k_1} + (1-\lambda)s_{k_2}, \cdots s_n) \\
&= \sum_{\mathbf{a} \in A} u_i(\mathbf{a}) \left(\prod_{j=1}^{k-1} s_j(a_j) \right) (\lambda s_{k_1}(a_k) + (1-\lambda)s_{k_2}(a_k)) \left(\prod_{j=k+1}^{n} s_j(a_j) \right) \\
&= \sum_{\mathbf{a} \in A} u_i(\mathbf{a}) \left(\prod_{j=1}^{k-1} s_j(a_j) \right) (\lambda s_{k_1}(a_k)) \left(\prod_{j=k+1}^{n} s_j(a_j) \right) \\
&\quad + \sum_{\mathbf{a} \in A} u_i(\mathbf{a}) \left(\prod_{j=1}^{k-1} s_j(a_j) \right) ((1-\lambda)s_{k_2}(a_k)) \left(\prod_{j=k+1}^{n} s_j(a_j) \right) \\
&= \lambda \sum_{\mathbf{a} \in A} u_i(\mathbf{a}) \left(\prod_{j=1}^{k-1} s_j(a_j) \right) s_{k_1}(a_k) \left(\prod_{j=k+1}^{n} s_j(a_j) \right) \\
&\quad + (1-\lambda) \sum_{\mathbf{a} \in A} u_i(\mathbf{a}) \left(\prod_{j=1}^{k-1} s_j(a_j) \right) s_{k_2}(a_k) \left(\prod_{j=k+1}^{n} s_j(a_j) \right) \\
&= \lambda U_i(s_1, \cdots, s_{k_1}, \cdots, s_n) + (1-\lambda)U_i(s_1, \cdots, s_{k_2}, \cdots, s_n)
\end{aligned}$$

The rest of the proof is straightforward as well. Let $s_{-i} \in S_{-i}$ and $b_1, b_2 \in B_i(s_{-i})$. Then $U_i(b_1) = U_i(b_2)$ and so

$$\begin{aligned}
U_i(\lambda b_1 + (1-\lambda)b_2, s_{-i}) &= \lambda U_i(b_1, s_{-i}) + (1-\lambda)U_i(b_2, s_{-i}) \\
&= \lambda U_i(b_1, s_{-i}) + (1-\lambda)U_i(b_1, s_{-i}) \\
&= U_i(b_1, s_{-i})
\end{aligned}$$

Then $\lambda b_1 + (1-\lambda)b_2 \in B_i(s_{-i})$, so $B_i(s_{-i})$ is convex for all $s_{-i} \in S_{-i}$ each player i. Thus, $B(s)$, the cartesian product of each $B_i(s_{-i})$ is also convex. □

And now we have all the mathematical machinery required to prove Nash's theorem for the existence of Nash equilibria for mixed extensions of finite n-person normal form games. The statement is as follows:

Theorem (Nash's Theorem)**.** *Every mixed extension of a finite n-person normal-form game has a Nash equilibrium*

Proof. Proof Let (N, S, U) be a mixed extension of a finite n-person normal-form game. It has been shown that S is a compact and convex set, and that the function $B : S \to S$ is upper-semi continuous. Moreover, $B(s)$ is convex for every $s \in S$. Then applying Kakutani's fixed point theorem, there exists $s_0 \in S$ such that $s_0 \in B(s_0)$. This is a Nash equilibrium as desired. □