Nash equilibrium and Sperner's lemma

Uriel Feige

May 7, 2024

1 Multiplayer Nash equilibrium

Recall that a (mixed) Nash equilibrium is a profile of mixed strategies, such that for every player, every strategy in the support of his mixed strategy is a best response (in expectation) to the profile of mixed strategies played by the other players.

Theorem 1.1 *Every finite game has a Nash equilibrium.*

A Nash equilibrium can be thought of as a "fixed point" of a game, in the sense that no player has an incentive to unilaterally deviate from it. (Note however that a player may be indifferent to deviating from a Nash equilibrium, in the sense of not caring how exactly to mix over his best response strategies.) Nash's proof of Theorem 1.1 was based on applying a different fixed point theorem, namely Brouwer's fixed point theorem that states that every continuous function from a compact convex body to itself has a fixed point.

Here is a sketch of how Nash's theorem can be proved using Brouwer's fixed point theorem. Let t be the number of players, and let n_i be the number of strategies available to player i. The set of all profiles of mixed strategies is a compact convex body G in R^d, where $d = \sum_{i=1}^{t} n_i$. One now defines a continuous function F from G to itself in a way that any fixpoint of F corresponds to a Nash equilibrium of the game. One such function is the following. Let $g_{i,j}(x)$ be the "gain" for player i with respect to profile x if he switches to pure strategy j. Observe that g is a continuous function (it is a polynomial in x). Define:

$$F(x)_{i,j} = \frac{x_{i,j} + \max[0, g_{i,j}(x)]}{1 + \sum_{k=1}^{n_i} \max[0, g_{i,k}(x)]}$$

Observe that in every point x of F, for every player i, at least for one value of j it must hold that $\max[0, g_{i,j}(x)] = 0$, because it could not be that every pure strategy is strictly better than the mixed strategy. Hence in a fixed point x of F, for all i and j it must hold that $\max[0, g_{i,j}(x)] = 0$. If j is in the support of the fixed point then it must be that $g_{i,j}(x) \geq 0$, as otherwise averaging implies that there must be some other strategy j' with $g_{i,j'}(x) > 0$. This means that the mixed strategy is composed only of best responses, and hence is a Nash equilibrium.

The reader may observe that the function F above can be thought of as a way of imposing dynamics on the space of profiles, where at a step one moves from profile x to profile $F(x)$. Though this dynamical system is guaranteed to have a fixed point, the proof does not guarantee that starting from an arbitrary point, a fixed point is ever reached (or that the dynamics eventually get arbitrarily close to a fixed point).

The proof of Nash's theorem is nonconstructive, due to the fact that Brouwer's fixed point theorem is nonconstructive. There is no polynomial time algorithm known for computing a Nash equilibrium. Part of the difficulty is that there are games (Nash shows such an example with three players) in which in every Nash equilibrium, the strategies in the support are played with nonrational probabilities (even though all payoffs have integer values). In a homework assignment we will see a three player game with an irrational Nash (though this game also has some rational Nash equilibria). Hence (at least in the case where there are more than two players), one would need to represent a Nash equilibrium either symbolically (for example, allow square root notation), or approximately using rational numbers. For standard models of computation, the latter option may be more convenient. Hence one may consider a notion of an ϵ-Nash, where $\epsilon > 0$ is a parameter that measures how close a profile of strategies is to being a true Nash. There are several notions of distance that one may consider (ϵ-close, ϵ-well supported), and for simplicity, we use the following one.

Normalize all payoffs so that they are between 0 and 1. Then one seeks a profile x in which for every player, his mixed strategy gives expected payoff within ϵ of the expected payoff of his best response.

2 Sperner's lemma in two dimensions

We present a combinatorial approach of Sperner for proving Brouwer's fixed point theorem. We sketch the proof for two dimensions, though it can be generalized to any dimension.

For simplicity of the presentation, assume that the compact convex set S in R^2 is an equilateral triangle, with vertices colored *red*, *blue*, and *green*. Consider a continuous function f with Lipschitz constant K from S to itself. If for some $x \in S$ it holds that $f(x) = x$, then we have a fixedpoint, and we are done. Color every non-fixpoint $x \in S$ according the direction to which f maps it, as follows. For a color c, if $f(x)$ is strictly closer than x to the edge facing the vertex of color c, then x receives color c. (For every non-fixpoint x, at least one color is eligible. If two colors are eligible, choose one of them arbitrarily.) Observe that verices receive their colors, and points on an edge between two vertices receive the color of one of the respective vertices. This is because f does not map points on the boundary of S to outside of S. Observe that if there are three points at distance ϵ from each other and with three different colors, two of them must be mapped by f to substantially different directions (an angle of at least $\pi/3$), and then the Lipschitz condition implies that they are mapped to a distance of $O(\epsilon K)$ from their current location. Letting ϵ tend to 0, this will give a fixpoint.

The situation can be abstracted as a Sperner coloring of a triangle. Subdivide the triangle by lines parallel to its three sides, with ϵ-distance between adjacent parallel lines. Each face in this subdivision is a triangle. Consider only the coloring of the intersection points of lines. We show that there is a multi-colored face (its three vertices have three distinct colors). In fact, for the purpose of applying induction to higher dimensions (which we will not do here), we prove that there is an odd number of multi-colored faces.

Given the colored subdivided triangle, we consider an auxiliary graph G. The vertices G are of two types – each 3-face is a vertex, and each boundary 2-face is a vertex. Two vertices are connected by an edge if their respective faces share a 2-face, with one endpoint of this 2-face colored red and the other colored blue. Hence depending on the coloring, vertices corresponding to boundary faces have degree either 0 or 1, and vertices corresponding to 3-dimensional faces have degree either 0, 1, or 2. Every graph has an even number of odd degree vertices. There is an odd number of odd degree vertices that correspond to boundary 2-faces (this fact which serves as a base case for a proof by induction over higher dimensions is easily proved, as the red-blue boundary has an odd number of color changes), and hence at least one vertex that corresponds to a 3-dimensional face has degree 1. The corresponding 3-face is necessarily 3-colored.