The Fundamental Theorem of Games

Yana Peysakhovich

September 23, 2007

1 Introduction

Two men are walking down the street. As they part company, police officers surround each man, arresting them separately, and they are brought in two separate cars to two separate police stations. Each man is seated in an interrogation room and immediately told the following by a detective:

"Look, we know it was you two who robbed the bank yesterday, but we want the charges to stick. We need you to cooperate with us. If you both confess, you'll get 5 years in prison. If you confess and your partner denies, you'll get 1 year, he'll get 10. If you both deny the crime, you'll both do 2 years. I'll be back in five minutes, and you better have a decision. Oh, and just so you know, your partner is getting the same offer."

The prisoners begin to get nervous.

What would you do? Maybe it depends on how much the prisoners trust one another, and maybe it depends on how much they hate jail. Well, if you are worried enough about jail, and don't care about what happens to your partner, you might make the perfect player. If you seek to maximize your gain and minimize your loss, if you are rational and intelligent, and if your rationality leads you to the conclusion that your 'opponent' is rational also, you might draw yourself this picture:

P1 \ P2	Confess	Deny
Confess	(-5,-5)	(-1,-10)
Deny	(-10,-1)	(-2,-2)

Then think: If the other prisoner confesses, I would be better off confessing ($-5 > -10$), and if he denies the crime, I would still be better off confessing ($-1 > -2$). So you confess. But the game is symmetric, so the other guy decides, similarly, to confess. You both confess and are sent to prison for 5 years. Too bad, certainly you would both have preferred the (Deny, Deny) outcome. So why could this outcome never be the result of this game, finitely played? Because if you know (or believe) that the other player will Deny, you will be very tempted to deviate from the (Deny, Deny) strategy. So would the other player. This makes you nervous. You decide to play it safe and confess. So does he. The strategy (Deny, Deny) is not *stable*, while the strategy (confess, confess) is stable. Neither player would prefer to change his response, given the strategy of the other player. In other words, at the point (Confess, Confess), player 1 says, given player 2 is playing "Confess," I would be only worse off playing "Deny," so I will stay where I am.

You have just played through a rather famous scenario in the field of game theory–The Prisoner's Dilemma, originally formulated by Merrill Flood and Melvin Dresher in 1950 [1]. Since then, Game Theory has come very far. This paper, however, will return to the basics, ultimately chronicling the evolution of the proof of the Fundamental Theorem of Games: every game has a solution.

2 The Beginning

One class of games is 'zero-sum.' These are simply 2-player games in which the sum of all payoffs between the players is 0, in particular, a gain of \$2 for one player is the loss of exactly \$2 for his opponent. The study of this class is quite fruitful, but limited in application to parlor games and extremely simplified international relations. However, it is precisely these games that gained the attention of the leading mathematicians and economists formulating the Theory of Games. Noting their definition, zero sum games can be intuitively represented in the form of an $m \times n$ matrix, where the payoff to payer one only is listed for every possible combination of strategies. Then, clearly, $0 - a_{ij}$ represents the payoff to player 2. Let's consider the game of Matching Pennies. Simultaneously, player 1 (P_1) and payer 2 (P_2) choose either "heads" or "tails." Each is in total ignorance of the other's decision. After the choices are made, they compare decisions, and P_2 pays P_1 \$1.00 if they match–if both chose heads or both chose tails. If they do not match, P_1 must pay P_2 \$1.00. But of course, this situation can easily be represented by the following simple 2×2 matrix:

$$\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

This is P_1's payoff matrix, where the first row and column both represent the choice of heads, while the second row and column represent the choice of tails. Entry a_{11} represents the outcome (Heads, Heads), entry a_{12} represents P_1 playing Heads and P_2 playing Tails. So, where P_1 has m possible moves and P_2 has n possible moves, we have the following matrix, in general:

$$\begin{pmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \vdots & \ddots & \vdots & \vdots \\ a_{m1} & a_{m2} & \dots & a_{mn} \end{pmatrix}$$

A few definitions will be useful:

Definition 2.1 *A* matrix game Γ *is given by any $m \times n$ matrix $A = (a_{ij})$ in which the entries a_{ij} are real numbers.*

Definition 2.2 *A* mixed strategy *for P_1 is an ordered m-tuple $\mathbf{x} = (x_1, x_2, ..., x_m)$ for non-negative real numbers x_i such that $x_1 + ... + x_m = 1$ Similarly, a* mixed strategy *for P_2 is an ordered n-tuple $\mathbf{y} = (y_1, y_2, ..., y_n)$ for non-negative real numbers y_j such that $y_1 + ... + y_n = 1$*

It is useful to note that in any game, players can only choose one strategy (ie, P_1 can only choose one of the m rows of the matrix (a_{ij}) and P_2 can only choose one of the n columns of the same matrix). A mixed strategy, however, is like creating a weighted die which, for P_1, has m sides and the weights on each side are dictated by the real numbers x_i. P_1 chooses a row by rolling his weighted die. Similarly, P_2's die has n sides, with the weight y_j corresponding to side j, and a roll of his die determines which column he chooses.

Definition 2.3 *A* pure strategy *for P_1 is the mixed strategy which is 1 in the i^{th} component and 0 everywhere else for $i = 1, ..., m$, and will be referred to as the i^{th}* pure strategy for P_1.

A similar definition holds for the j^{th} *pure strategy of* P_2.

Definition 2.4 *The* payoff function *for Γ is defined to be*

$$E(\mathbf{x}, \mathbf{y}) = \sum_{i,j} x_i a_{ij} y_j$$

where $\mathbf{x} = (x_1, x_2, ..., x_m)$ and $\mathbf{y} = (y_1, y_2, ..., y_n)$ are mixed strategies.

Definition 2.5 *A* solution *of Γ is a pair of mixed strategies $\overline{\mathbf{x}} = (\overline{x}_1, \overline{x}_2, ..., \overline{x}_m), \overline{\mathbf{y}} = (\overline{y}_1, \overline{y}_2, ..., \overline{y}_n)$, and a real number v such that $E(\overline{\mathbf{x}}, j) \geq v$ for the pure strategies $j = 1, ..., n$ and $E(i, \overline{\mathbf{y}}) \leq v$ for the pure strategies $i = 1, ..., m$. Then $\overline{\mathbf{x}}$ and $\overline{\mathbf{y}}$ are called* optimal strategies *and the number v is called the* value *of the game Γ.*

In fact, the amount v (the value of the game) is the expectation that P_1 can assure himself *no matter what* P_2 does, while P_2 can protect himself against expectations higher than v *no matter what* P_1 does.

At last, we can begin our investigation of a remarkable fact which was proven by John von Neumann (arguably, the father of Game Theory) in the late 1920's: Every matrix game has a solution[2].

3 The Fundamentals of Matrix Games

Before we prove this fact, however, we will briefly investigate the case for a 2×2 zero-sum game (matrix game)

3.1 The 2x2 Case

Theorem 1 *Given the* 2×2 *matrix*

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$$

there exist vectors $\overline{\mathbf{x}} = (\overline{x}_1, \overline{x}_2)$*, where* $\overline{x}_1, \overline{x}_2 \geq 0$ *and* $\overline{x}_1 + \overline{x}_2 = 1$*, and* $\overline{\mathbf{y}} = (\overline{y}_1, \overline{y}_2)$*, where* $\overline{y}_1, \overline{y}_2 \geq 0$ *and* $\overline{y}_1 + \overline{y}_2 = 1$*, and a real number* v *such that:*

(1)

$\overline{x}_1 a_{11} + \overline{x}_2 a_{21} \geq v$

$\overline{x}_1 a_{12} + \overline{x}_2 a_{22} \geq v$

(2)

$a_{11}\overline{y}_1 + a_{12}\overline{y}_2 \leq v$

$a_{21}\overline{y}_1 + a_{22}\overline{y}_2 \leq v$

Proof 1 If P_2 plays the mixed strategy $\mathbf{y} = (y_1, y_2)$, he can expect to pay $E(1, \mathbf{y}) = a_{11}y_1 + a_{12}y_2$ against 1 by P_1 and $E(2, \mathbf{y}) = a_{21}y_1 + a_{22}y_2$ against 2 by P_1. Plotting $(E(1, \mathbf{y}), E(2, \mathbf{y}))$ in the case where $y_1 = 1$ and $y_2 = 0$, then in the case where $y_1 = 0$, $y_2 = 1$, we obtain the points (a_{11}, a_{21}) and (a_{12}, a_{22}) both in $\mathbf{R^2}$. The line connecting these points creates P_2's *expectation space*, since this gives all expectations under mixed strategies (where $y_1, y_2 \geq 0$, $y_1 + y_2 = 1$). We can notice that P_2 is assured of paying no more than the expectation $e \Leftrightarrow$ his expectation point lies in the set S_e of all points with both coordinates less than or equal to e. Since P_2 is playing optimally when we choose a strategy $\mathbf{y}$ with the smallest possible e, we translate S_e until it is just in contact with P_2's expectation space. We will call this translated set S_v. It is clear that condition (2) of the theorem is satisfied.

Suppose we have found a mixed strategy $\overline{\mathbf{x}} = (\overline{x}_1, \overline{x}_2)$ satisfying (1) for this v. Then, the line $\overline{x}_1 t_1 + \overline{x}_2 t_2 = v$, *separates* P_2's expectation space from the the set S_v, because we have $\overline{x}_1 t_1 + \overline{x}_2 t_2 \geq v$ for all points (t_1, t_2) in P_2's expectation space by (1) and we have $\overline{x}_1 t_1 + \overline{x}_2 t_2 \leq \overline{x}_1 v + \overline{x}_2 v = v \ \forall$ points $(t_1, t_2) \subseteq S_v$ (by definition).

Now we have found that any optimal $\mathbf{x}$, if it exists, defines a separating line. Let us assume that $at_1 + bt_2 = c$ separates P_2's expectation space from S_v, that is,

(3) $aa_{11} + ba_{12} \geq c$
(4) $aa_{12} + ba_{22} \geq c$
(5) $at_1 + bt_2 \leq c \ \forall (t_1, t_2) \subseteq S_v$.
Since $(v-1, v)$ and $(v, v-1)$ are in S_v, we have
(6) $a(v-1) + bv \leq c$
(7) $av + b(v-1) \leq c$
and, because (v, v) lies on any separating line,
(8) $av + bv = c$

By subtracting (8) from (6) we see that $a \geq 0$, and by subtracting (8) from (7) we see that $b \geq 0$. But a and b cannot both be 0, since they define a separating line, so we define $\overline{x}_1 = \frac{a}{a+b}$, $\overline{x}_2 = \frac{b}{a+b}$, so $\overline{x}_1, \overline{x}_2 \geq 0$, $\overline{x}_1 + \overline{x}_2 = 1$. Then, (3) and (4) give

$\overline{x}_1 a_{11} + \overline{x}_2 a_{21} \geq \frac{c}{a+b}$ and
$\overline{x}_1 a_{12} + \overline{x}_2 a_{22} \geq \frac{c}{a+b}$,
which, combined with (8) give the desired inequalities (1).

3.2 Convex Sets

The remainder of this section will assume basic knowledge of linear algebra. We will be working in the Euclidean vector space $\mathbf{R^m}$. The *norm* in our space will be characterized by $\|\mathbf{t}\| = \sqrt{\mathbf{t} \cdot \mathbf{t}}$, Where the 'dot' product $\mathbf{t} \cdot \mathbf{u} = \sum_{i=1}^{m} t_i u_i$.

Definition 3.1 *A* hyperplane $H(\mathbf{x}, a)$ *in* $\mathbf{R^m}$ *is the set of all vectors* $\mathbf{t}$ *such that* $\mathbf{x} \cdot \mathbf{t} = a$ *for a given vector* $\mathbf{x} \neq 0$ *and a real number* a.

Naturally, we can refer to
$H^+(\mathbf{x}, a) = \{\mathbf{t} | \mathbf{x} \cdot \mathbf{t} \geq a\}$
$H^-(\mathbf{x}, a) = \{\mathbf{t} | \mathbf{x} \cdot \mathbf{t} \leq a\}$, as the upper and lower hyperplane, respectively.

Definition 3.2 *A subset C of* $\mathbf{R^n}$ *is called* convex *if whenever* $\mathbf{t_1}, ..., \mathbf{t_m} \in C$ *and* $b_1, ..., b_m$ *are non-negative real numbers summing to* 1, $\mathbf{t} = b_1\mathbf{t_1} + ... + b_m\mathbf{t_m} \in C$. *The vector* $\mathbf{t}$ *is called a* convex combination *of* $\mathbf{t_1}, ..., \mathbf{t_m}$.

We can notice that, geometrically, a convex set contains all line segments joining two points in the set.

Definition 3.3 *The* convex hull $C(S)$ *of a given set S is the set of all convex combinations of sets of points from S, or, it is the minimal convex set containing S.*

Theorem 2 *Given a closed convex set $C \subset \mathbf{R^n}$ and a vector $\mathbf{u} \notin C$, there exists a hyperplane H such that:*
1) $C \subset H^+$
2) $\mathbf{u} \in H^-$ but $\mathbf{u} \notin H$.

Proof 2 Because we have a distance function, $\|\mathbf{t}-\mathbf{u}\|$, which is continuous on a closed and bounded set of $\mathbf{t} \in C$ for which $\|\mathbf{t}-\mathbf{u}\| \leq \|\mathbf{t_0}-\mathbf{u}\|$ for some fixed $\mathbf{t_0}$ in C, the distance attains its minimum in C, so we can choose the point $\mathbf{v} \in C$ which is closest to $\mathbf{u}$. To simplify our notation, we will assume $\mathbf{v} = \mathbf{0}$.

We now consider the hyperplane $H(-\mathbf{u}, 0)$. It is clear that condition (2) of the theorem is met, since $-\mathbf{u}\cdot\mathbf{u} < 0$ (we know $\mathbf{u} \neq \mathbf{0}$).

Suppose that (1) does not hold. Then there is a vector $\mathbf{t} \in C$ for which $-\mathbf{u}\cdot\mathbf{t} < 0$, so for which $\mathbf{u}\cdot\mathbf{t} > 0$. Certainly, either $\mathbf{u}\cdot\mathbf{t} \geq \mathbf{t}\cdot\mathbf{t}$, or $0 < \mathbf{u}\cdot\mathbf{t} < \mathbf{t}\cdot\mathbf{t}$. We will call this Case 1 and Case 2, respectively.

Case 1:

$$\mathbf{u}\cdot\mathbf{t} \geq \mathbf{t}\cdot\mathbf{t} \Rightarrow \|\mathbf{t}-\mathbf{u}\|^2 = \mathbf{t}\cdot\mathbf{t} - 2\mathbf{u}\cdot\mathbf{t} + \mathbf{u}\cdot\mathbf{u} = \mathbf{u}\cdot\mathbf{u} - (\mathbf{u}\cdot\mathbf{t} - \mathbf{t}\cdot\mathbf{t}) - \mathbf{u}\cdot\mathbf{t} < \mathbf{u}\cdot\mathbf{u} = \|\mathbf{u}\|^2 \Rightarrow \|\mathbf{t}-\mathbf{u}\| < \|\mathbf{u}\| = \|\mathbf{0}-\mathbf{u}\|,$$

But this contradicts the fact that $\mathbf{0}$ is the point of C closest to $\mathbf{u}$.

Case 2:

$$0 < \tfrac{\mathbf{u}\cdot\mathbf{t}}{\mathbf{t}\cdot\mathbf{t}} < 1 \Rightarrow (\tfrac{\mathbf{u}\cdot\mathbf{t}}{\mathbf{t}\cdot\mathbf{t}})\mathbf{t} \in C, \text{ and}$$
$$\|(\tfrac{\mathbf{u}\cdot\mathbf{t}}{\mathbf{t}\cdot\mathbf{t}})\mathbf{t} - \mathbf{u}\|^2 = \mathbf{u}\cdot\mathbf{u} - (\tfrac{\mathbf{u}\cdot\mathbf{t}}{\mathbf{t}\cdot\mathbf{t}})(\mathbf{u}\cdot\mathbf{t}) < \|\mathbf{u}\|^2 \Rightarrow \|(\tfrac{\mathbf{u}\cdot\mathbf{t}}{\mathbf{t}\cdot\mathbf{t}})\mathbf{t} - \mathbf{u}\| < \|\mathbf{u}\|$$

So $(\frac{\mathbf{u}\cdot\mathbf{t}}{\mathbf{t}\cdot\mathbf{t}})\mathbf{t} \in C$, and again is closer to $\mathbf{u}$ than 0, a contradiction.

Both possible cases have led to a contradiction, thus our assumption is incorrect and (1) holds.

Definition 3.4 *If H is a hyperplane and C is a convex set such that C is contained in H^+, then H is called a* supporting hyperplane *for C and H^+ is called a* support *for C.*

Theorem 3 *Any closed convex set is the intersection of its supports.*

Proof 3 Let $\bigcap H^+$ denote the intersection of all the supports for C. Clearly, $C \subseteq \bigcap H^+$. However, for any $\mathbf{u} \notin C$, there is a support H^+ for C which does not contain $\mathbf{u}$, by Theorem 2. Thus, $C \supseteq \bigcap H^+$ and so we see that $C = \bigcap H^+$.

Corollary 1 *Given a closed convex set C such that $C \neq \mathbf{R^n}$, there exists a support for C which contains a vector in C.*

This corollary follows immediately from Theorem 2.

Definition 3.5 *The point* $\mathbf{u}$ *is an* interior point *of a set* $S \subseteq \mathbf{R^n}$ *if there exists* $\varepsilon > 0$ *such that if* $\|\mathbf{v} - \mathbf{u}\| < \varepsilon$ *then* $\mathbf{v} \in S$

Theorem 4 *Given a set* $S \subset \mathbf{R^n}$ *consisting of the vector* $\mathbf{0}$ *and* n *LI vectors* $\mathbf{t_1}, ..., \mathbf{t_n}$*,* $C(S)$ *contains an interior point.*

Proof 4 We can note that in fact $C(S)$ is the intersection of its extreme supports–those which contain n points from the set S. We simply modify Theorem 3 to see that the result holds.

Theorem 5 *Every convex set* $C \subseteq \mathbf{R^n}$ *either contains an interior point or is contained in a hyperplane.*

Proof 5 We can assume that $\mathbf{0} \in C$, for simplicity. Then, we can chose a system of LI vectors $\mathbf{t_1}, ..., \mathbf{t_m}$ from C such that every set of $m+1$ vectors from C is LD. Then, either:

Case 1:

If $m = n$, then we are done by the previous theorem.

Case 2:

If $m < n$, then the subspace $\mathbf{R^m}$ generated by $\mathbf{t_1}, ..., \mathbf{t_m}$ is $\leq \mathbf{R^n}$ and we can choose a non-zero vector $\mathbf{x}$ such that $\mathbf{x} \cdot \mathbf{t_k} = 0 \ \forall \ k = 1, ..., m$. But, for any $\mathbf{t} \in C$, we have $\mathbf{t} = \alpha_1 \mathbf{t_1} + ... + \alpha_m \mathbf{t_m}$, since any set of $m+1$ vectors from C is LD, while $\mathbf{t_1}, ..., \mathbf{t_m}$ is not. Then $\mathbf{x} \cdot \mathbf{t} = 0 \Rightarrow C \subset H(\mathbf{x}, 0)$.

Theorem 6 *Given any convex set* C *with* $C \neq \mathbf{R^n}$*, there exists a support for* C*.*

Proof 6 Consider D, the smallest closed set containing C, that is, $D := \overline{C}$. Omitting details, it is easy to show that if there exists a sequence $\mathbf{u_1}, \mathbf{u_2}, \mathbf{u_3}, ... \longrightarrow \mathbf{u}$ and a sequence $\mathbf{v_1}, \mathbf{v_2}, \mathbf{v_3}, ... \longrightarrow \mathbf{v}$, so that $\mathbf{u}$ and $\mathbf{v}$ are limit points, as $\mathbf{u_k} \longrightarrow \mathbf{u}$ and $\mathbf{v_k} \longrightarrow \mathbf{v}$, $a\mathbf{u_k} + b\mathbf{v_k} \longrightarrow a\mathbf{u} + b\mathbf{v}$, for $a, b \in \mathbf{R}$. So D is convex. By Theorem 2, we only need to see that D is not all of $\mathbf{R^n}$. By Theorem 5, however, we can note that if $C \subseteq H$, a hyperplane, then D must also be in H. If C has an interior point $\mathbf{u}$ and there exists $\mathbf{v} \notin C$, then for some $\varepsilon > 0$, all points $\mathbf{x} \in B(2\mathbf{v} - \mathbf{u}, \varepsilon) \notin C$, thus, $2\mathbf{v} - \mathbf{u} \notin D$.

At this time we can note an intuitive definition, given what we have seen so far:

Definition 3.6 *The Hyperplane* $H(\mathbf{x}, a)$ *is said to* separate *the set* C *from the set* D *if* $C \subset H^+$ *and* $D \subset H^-$*.*

Theorem 7 *Let two convex sets C and D be such that:*

(1) $\mathbf{0} \in C, D$;

(2) D has an interior point $\mathbf{w}$;

(3) no point of C is interior to D.

Then there exists $H(\mathbf{x}, 0)$ which separates C from D.

Proof 7 Let $E = \{a\mathbf{t} - b\mathbf{u} \mid a, b \geq 0, \mathbf{t} \in C, \mathbf{u} \in D\}$. E is, by definition, a convex set. We can also notice that any supporting $H(\mathbf{x}, 0)$ for E separates C from D. Finally, we must see that $\mathbf{w} \notin E$.

Suppose $\mathbf{w} = a\mathbf{t} - b\mathbf{u}, a, b \geq 0, \mathbf{t} \in C, \mathbf{u} \in D$. Then

$$\tfrac{1}{1+b}\mathbf{w} + \tfrac{b}{1+b}\mathbf{u} = \tfrac{a}{1+b}\mathbf{t}.$$

If $\frac{a}{1+b} \leq 1$, then the R.H.S. of the equation above is a convex combination of $\mathbf{0}$ and $\mathbf{t}$, hence is in C, while the L.H.S. is a convex combination of $\mathbf{u} \in D$ and $\mathbf{w}$, which is interior to D, hence the L.H.S. itself is interior to D. But this is a contradiction to the fact that no point of C is interior to D.

If $\frac{a}{1+b} \geq 1$, then, since $\mathbf{t} \in C$, and in this case, is also a convex combination of $\mathbf{0}$ and $\frac{a}{1+b}\mathbf{t}$, which is interior to D by the previous argument. So we have that $\mathbf{t}$ is also interior to D. Obviously, this is a contradiction. So, $\mathbf{w} \notin E$.

Then, since E is a convex set, $\mathbf{w} \notin E$, and by Theorem 6, $\exists$ a supporting hyperplane $H(\mathbf{x}, a)$ for E. But then it must be true that $H(\mathbf{x}, 0)$ also supports E:

Suppose not, then $\mathbf{x} \cdot \mathbf{v} < 0$ and $\mathbf{x} \cdot \mathbf{v} \geq a$ for some $\mathbf{v}$ in E. Then for some $c > 0$, $c(\mathbf{x} \cdot \mathbf{v}) + \mathbf{x} \cdot \mathbf{v} < a$. Then $\mathbf{x} \cdot (c\mathbf{v} + \mathbf{v}) < a$ and $c\mathbf{v} + \mathbf{v} = (c+1)\mathbf{v} \in E$, which is a contradiction.

Thus, $H(\mathbf{x}, 0)$ separates C from D

3.3 The Fundamental Theorem for Matrix Games

Theorem 8 *Every matrix game has a solution. In other words:*

Given any $m \times n$ matrix $A = (a_{ij})$, there exist vectors

$\overline{x} = (\overline{x}_1, ..., \overline{x}_m)$, $\overline{x}_i \geq 0 \ \forall \ i$, $\overline{x}_1 + \cdots + \overline{x}_m = 1$,

$\overline{y} = (\overline{y}_1, ..., \overline{y}_n)$, $\overline{y}_j \geq 0 \ \forall \ j$, $\overline{y}_1 + \cdots + \overline{y}_n = 1$

and a real number v s.t.

(1)

$\overline{x}_1 a_{1j} + \cdots + \overline{x}_m a_{mj} \geq v$ *for* $j = 1, ..., n$

(2)

$a_{i1}\overline{y}_1 + \cdots + a_{in}\overline{y}_n \leq v$ *for* $i = 1...., m$

Proof 8 Again, we will consider P_2's expectation space in $\mathbf{R^m}$ and consider the set C of all vectors $\mathbf{t} = y_1\mathbf{t_1} + \cdot + y_n\mathbf{t_n}$ where $\mathbf{t_j} = (a_{1j}, ..., a_{mj})$, all $y_i \geq 0$, $\overline{y}_1 + \cdots + \overline{y}_n = 1$.

This set C is the convex hull of the points $\mathbf{t_1}, ..., \mathbf{t_n}$. If $\mathbf{z} = (z_1, ..., z_m)$ is a point in C, we can consider the function $\psi(\mathbf{z}) = max_{i=1,...,m}\{z_i\}$. $\psi(\mathbf{z})$ is a continuous function defined on the compact set $C \subset \mathbf{R^m}$ and hence attains its minimum at some point $\overline{\mathbf{t}} = \overline{y}_1\mathbf{t_1} + \cdots + \overline{y}_n\mathbf{t_n} \in C$. Let $v = \psi(\overline{\mathbf{t}})$. It is clear that (2) is satisfied for this choice of $\overline{\mathbf{y}}$ and v.

Now we define D to be the set of all $\mathbf{t} = (t_1, ..., t_m)$ such that $t_i \leq v$ for $i = 1, ..., m$. The set D intersects C at $\overline{\mathbf{t}}$ and has interior points (clearly). However, no point of C is interior to D by our choice of v. But these are exactly the conditions we need to employ Theorem 7, so we know $\exists$ a hyperplane $H(\mathbf{x}, a)$ which separates C from D.

In fact, the point $\mathbf{v} = (v, ..., v)$ lies in H. To see that this is true, suppose not, then

$$\mathbf{x} \cdot \mathbf{v} < a \Longrightarrow \mathbf{x} \cdot (2\overline{\mathbf{t}} - \mathbf{v}) = 2\mathbf{x} \cdot \overline{\mathbf{t}} - \mathbf{x} \cdot \mathbf{v} < 2a - a = a$$

This contradicts the fact that $2\overline{\mathbf{t}} - \mathbf{v} \in D$. It is true from the definition of D that $\mathbf{x} \cdot \mathbf{v} \not> a$. Now we have that $\mathbf{x} \cdot \mathbf{v} = a$ and $\mathbf{v} \in H$.

We let $\mathbf{v_i}$ be the vector which is v in every component except at the i^{th} coordinate, where it is $v - 1$, for $i = 1, ..., m$. Then $\mathbf{v_i} \in D$ and so $\mathbf{x} \cdot \mathbf{v_i} = (\mathbf{x} \cdot \mathbf{v}) - x_i \leq a$. Therefore, $x_i \geq 0$ for $i = 1, ..., m$. Since $\mathbf{x} \neq \mathbf{0}$, we can define $\overline{\mathbf{x}} = (\overline{x}_1, ..., \overline{x}_m)$ where

$$\overline{x}_i = \tfrac{x_i}{x_1 + \cdots + x_m} \text{ for } i = 1, ..., m.$$

Then all $\overline{x}_i \geq 0$, $\overline{x}_1 + \cdots + \overline{x}_m = 1$. Since $C \subset H^+$,

$$\overline{\mathbf{x}} \cdot \mathbf{t_j} = \tfrac{1}{x_1 + \cdots + x_m}\mathbf{x} \cdot \mathbf{t_j} \geq \tfrac{a}{x_1 + \cdots + x_m} = v$$

for $j = 1, ..., n$, which is condition (1) of the Theorem.

4 A Fixed Point Theorem

Before we are able to prove John Nash's remarkable generalization of the previous result, we must familiarize ourselves with a new topic. We will first prove two lemmas, and from these the *Brouwer Fixed Point Theorem* will follow.

Definition 4.1 *Let* $\mathbf{x_1}, \mathbf{x_2}, ..., \mathbf{x_{n+1}}$ *be* $n + 1$ *points in* $\mathbf{R^k}$ *where* $k \geq n$ *and any* n *of the points are linearly independent. Then the* n-simplex *defined by*

$\mathbf{x_1}, \mathbf{x_2}, ..., \mathbf{x_{n+1}}$ *is the set* S *of convex combinations of* $\mathbf{x_1}, \mathbf{x_2}, ..., \mathbf{x_{n+1}}$*:*

$$S = \{\mathbf{x} \mid \mathbf{x} = \sum_{i=1}^{n+1} \lambda_i \mathbf{x_i}, \lambda_i \geq 0, \sum_{i=1}^{n+1} \lambda_i = 1\}$$

Definition 4.2 *For* $\mathbf{x} \in S$*, we define* λ_i *(as above) to be the* i^{th} barycentric coordinate *of* $\mathbf{x}$*.*

Definition 4.3 *The points* $\mathbf{x_1}, \mathbf{x_2}, ..., \mathbf{x_{n+1}}$ *are the* vertices *of* S*. We will label the vertex* x_i *with the number* i*.*

Definition 4.4 *For a given* $\mathbf{x} \in S$*, the set* $\{\mathbf{x_i} \mid \lambda_i > 0\}$ *will be called the* carrier *of* $\mathbf{x}$*.*

Definition 4.5 *A* face*,* F*, of the simplex* S *is defined as*

$$F = \{\mathbf{x} \mid \mathbf{x} = \textstyle\sum_{i=1}^{n+1} \lambda_i \mathbf{x_i}, \lambda_k = 0 \text{ for one } k, \lambda_i \geq 0, \sum_{i=1}^{n+1} \lambda_i = 1\}.$$

Definition 4.6 *A triangulation, or,* simplicial subdivision *of* S *is a finite family of* $\{S_j\}$ *so that*
(i) *The elements of* $\{S_j\}$ *have disjoint interiors*
(ii) *If a vertex* $\mathbf{x_i}$ *of* S_j *is an element of* S_k*, then* $\mathbf{x_i}$ *is also a vertex of* S_k*.*
(iii) $\bigcup S_k = S$*.*

Definition 4.7 *Let* $\{S_j\}$ *be a simplicial subdivision of* S*. We* label *each vertex of each subsimplex with one of the numbers* $1, 2, ..., n+1$*. A labeling is said to be* admissible *if each vertex is labeled with the index of one of the elements of its carrier.*

Theorem 9 (Sperner's Lemma) *Let* $\{S_j\}$ *be a simplicial subdivision of the n-simplex* S*. If* $\{S_j\}$ *is labeled admissibly, then there exists* $S_0 \in \{S_j\}$ *such that* S_0 *is completely labeled, i.e. with labels of each number* $1, 2, ..., n+1$*.*

Proof 9 We will, in fact, prove that the number of completely labeled subsimplices is odd. We proceed by induction.

Consider n=1: Each vertex of our simplicial subdivision is labeled 1 or 2 (our simplex is a line segment). Let a be the number of subsegments whose end points are both labeled 1 and let b be the number of subsegments whose end points are completely labeled, so one end point is labeled 1 and the other is labeled 2. We will now count the number of endpoints labeled 1, where each endpoint is counted once for each subsegment of which it is an element. Obviously, this number is $2a + b$. Now, we will count again. We have one "exterior" endpoint labeled 1, the endpoint of the original line segment, and, if we let c be the number of interior endpoints labeled 1, we see that the total number of endpoints labeled 1 is $1 + 2c$. So, $1 + 2c = 2a + b$. Since $1 + 2c$ is odd, b must be odd.

For the remainder of the proof, it may be useful to refer to the figure below, an admissibly labeled simplicial subdivision of a simplex.

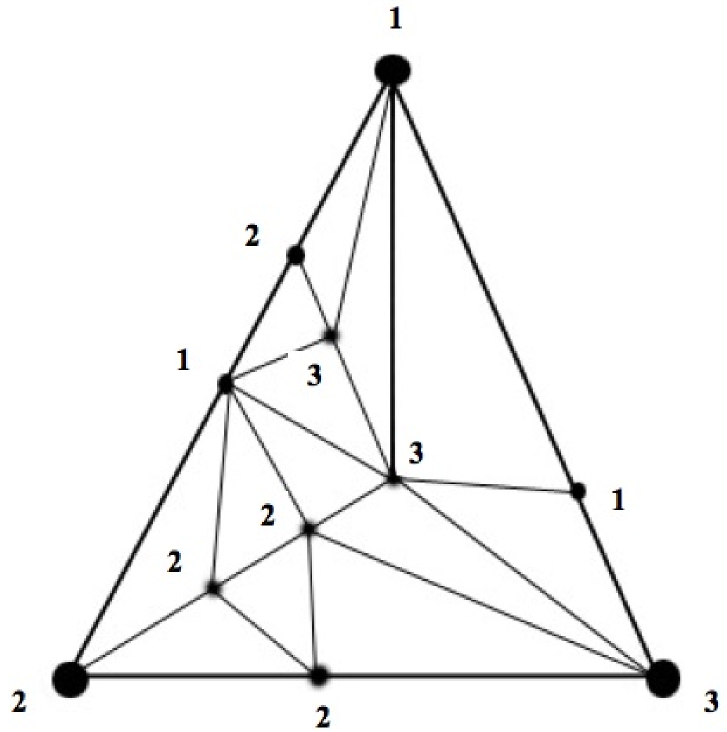

Now, suppose for an $(n-1)$-simplex, any admissibly labeled simplicial subdivision contains an odd number of completely labeled subsimplices.

Consider an admissibly labeled subdivision of an n-simplex. An argument similar to the case for $n = 1$ holds here. We let a be the number of subsimplices labeled with $1, ..., n$ but not $n + 1$. For each of these subsimplices, there are two faces with the labels $1, ..., n$. Let b be the number of completely labeled subsimplices. Each of these completely labeled subsimplices has exactly one face with the labels $1, ..., n$. So the total number of faces with the labels $1, ..., n = 2a + b$. Now we count again. Let c be the number of interior faces carrying the labels $1, ..., n$. Again, each interior face must be a face of precisely 2 adjacent subsimplices. Exterior faces can be defined as those not shared by two adjacent subsimplices, so clearly, each exterior face labeled $1, ..., n$ will be counted once. The number of such faces, however, must be odd, by the inductive hypothesis, since a face of an n-simplex is an

$(n-1)$-simplex. Thus, if we let d be the number of exterior faces carrying the labels $1, ..., n$, we see that $2a + b = 2c + d$. Because d is odd, b is odd.

Thus, by induction, the number of completely labeled subsimplices is odd. Since 0 is not odd, the theorem is proven

Theorem 10 (Knaster-Kuratowski-Mazurkewicz (KKM) Theorem) *Let S be an n-simplex. Consider the sets $C_1, C_2, ..., C_{n+1} \subset S$ where C_j is closed, and where vertex $j = \mathbf{x_j} \in C_j$. For all $\mathbf{x}$ in S, let $\mathbf{x} \in C_i$ for some i s.t. $\mathbf{x_i}$ is a carrier of $\mathbf{x}$. Then*

$$\cap_{j=1}^{n+1} C_j \neq \emptyset$$

Proof 10 We can choose a sequence of simplicial subdivisions $\Lambda_v = \{S_v^k \mid k = 1, 2, ...\}$ for $v = 1, 2, ...$ Then k indexes the subsimplices within each subdivision Λ_v. We construct the sequence $\Lambda_1, \Lambda_2, \Lambda_3, ...$ such that for $v < u$, the minimal ε-Ball which contains S_u^k has smaller radius that the minimal ε-Ball which contains S_v^j for all reasonable values of j, k. In other words, the 'mesh' of our sequence Λ_v will become progressively finer and arbitrarily fine as v increases. We can now label the vertices of S_v^k by the number j, where the vertex is an element of C_j for some j s.t. $\mathbf{x_j}$ is an element of the carrier of the vertex.

By Sperner's Lemma, we know that there exists some $S_v^0 \in \Lambda_v$, so that S_v^0 is completely labeled. Let $\mathbf{x}_v^i$ be the vertex of S_v^0 with the label i. Then $\mathbf{x}_v^i \in C_i$ for all v. But S is compact, so there exists a convergent subsequence of $\{\mathbf{x}_v^i\}_{v=1}^{v=\infty}$ for each i. But because the diameter of our subsimplex goes to zero, these subsequences must converge to a common point, $\mathbf{x}_0$, for all i. Since C_i is closed, $\mathbf{x}_v^i \longrightarrow \mathbf{x}_0 \Longrightarrow \mathbf{x}_0 \in C_i$ for all i. Thus $\mathbf{x}_0 \in \cap_{j=1}^{n+1} C_j \neq \emptyset$

[3]

Theorem 11 (Brouwer Fixed-Point Theorem) *Let S be an n-simplex and let $f : S \longrightarrow S$, f continuous. Then there exists $\mathbf{x}^* \in S$ s.t. $f(\mathbf{x}^*) = \mathbf{x}^*$*

Proof 11 Let $\lambda_j(\mathbf{x})$ be the j^{th} barycentric coordinate of $\mathbf{x}$. Define

$$C_j = \{x \mid \lambda_j(f(\mathbf{x})) \leq \lambda_j(\mathbf{x})\}$$

In fact, C_j satisfies the assumptions of the KKM theorem above by the continuity of f (this is easy to check).

Then, applying the KKM theorem to these C_j's, there exists $\mathbf{x}^* \in S$ s.t. $\mathbf{x}^* \in \cap_{j=1}^{n+1} C_j$. We have, then, by our definition, that

$$\lambda_j(f(\mathbf{x}^*)) \leq \lambda_j(\mathbf{x}^*) \text{ for all } j.$$

However, since $\sum \lambda_j(\mathbf{x}^*) = \sum \lambda_j(f(\mathbf{x}^*)) = 1$, it must be that $\lambda_j(\mathbf{x}^*) = \lambda_j(f(\mathbf{x}^*))$ for all j, or, in other words, $\mathbf{x}^* = f(\mathbf{x}^*)$.

[3]

5 The Triviality

It is reported that when young John Nash came to von Neumann with his proof of the Fundamental Theorem for Non-Zero Sum games, the envious von Neumann dismissed his result with the phrase, "but it's a triviality" [5]. So it may seem to us today. Nonetheless, Nash accomplished what no one else had, and his contribution to the field of Game Theory certainly rivals that of von Neumann.

We will adopt John Nash's notation for elements of games [6]:

Definition 5.1 *An n-*person game *is a set of n players, or, 'positions' each with an associated finite set of* pure strategies*, and for each player i, there is a payoff function $p_i : \{n\text{-tuples of pure strategies}\} \longrightarrow \mathbf{R}$*

Definition 5.2 *A* mixed strategy *for player i will be a collection of non-negative numbers c_{i_α} so that*

$$\mathbf{s}_i = \sum_\alpha c_{i_\alpha} \pi_{i_\alpha}, c_i \geq 0, \sum_\alpha c_{i_\alpha} = 1$$

where $\mathbf{s}_i$ is the mixed strategy of player i and π_{i_α} is the α^{th} pure strategy of player i

We can think of the $\mathbf{s}_i$'s as points in a simplex with vertices π_{i_α}. This simplex is a convex subset of a real vector space.

The symbols i, j, k will refer to players i, j, k

The symbols α, β, γ will indicate various pure strategies of a player

The symbols $\mathbf{s}_i, \mathbf{t}_i, \mathbf{r}_i$ will indicate mixed strategies

We now extend the definition of p_i above to the function $p_i : \{n$-tuples of mixed strategies$\} \longrightarrow \mathbf{R}$, so that $p_i(\mathbf{s}_1, \mathbf{s}_2, ..., \mathbf{s}_n)$ will represent the payoff to player i under the mixed strategies $\mathbf{s}_1$ for player 1, $\mathbf{s}_2$ for player $2, ..., \mathbf{s}_n$ for player n.

And of course, as in the zero-sum version, a pure strategy α for player $i = \pi_{i_\alpha}$is simply a mixed strategy $\mathbf{s}_i = 0\pi_{i_1} + 0\pi_{i_2} + \cdots + 1\pi_{i_\alpha} + 0\pi_{i_{\alpha+1}} + \cdots + 0\pi_{i_n}$.

We will want to be sure that any mixed strategy is *stable*, in order to find an equilibrium, so it will be helpful to introduce the following notation:

Where $\mathcal{S} = (\mathbf{s}_1, \mathbf{s}_2, ..., \mathbf{s}_n)$, substitutions will be represented by $(\mathcal{S}; \mathbf{t}_i) = (\mathbf{s}_1, \mathbf{s}_2, ..., \mathbf{s}_{i-1}, \mathbf{t}_i, \mathbf{s}_{i+1}, ..., \mathbf{s}_n)$. Successive substitutions, $((\mathcal{S}; \mathbf{t}_i); \mathbf{r}_j)$ will be indicated by $(\mathcal{S}; \mathbf{t}_i; \mathbf{r}_j)$, etc.

Definition 5.3 *An n-tuple $\mathcal{S}$ is an equilibrium point if and only if:*

$$\forall i, p_i(\mathcal{S}) = \max_{all \mathbf{r}_i' s} p_i(\mathcal{S}; \mathbf{r}_i)$$

So that in equilibrium, each player's strategy is optimal against all others'.

Definition 5.4 *A mixed strategy* $\mathbf{s}_i$ uses *a pure strategy* π_{i_α} *if* $\mathbf{s}_i = \sum_\beta c_{i_\beta}\pi_{i_\beta}$ *and* $c_i > 0$. *If* $\mathcal{S} = (\mathbf{s_1}, ..., \mathbf{s_n})$ *and* $\mathbf{s_i}$ *uses* π_{i_α}*, we say that* $\mathcal{S}$ *also uses* π_{i_α}

Definition 5.5 *From the linearity of* $p_i(\mathbf{s_1}, ..., \mathbf{s_n})$ *in* $\mathbf{s_i}$*,*

$$\max_{all\mathbf{r}_i's} p_i(\mathcal{S}; \mathbf{r}_i) = \max_\alpha p_i(\mathcal{S}; \pi_{i_\alpha})$$

We define $p_{i_\alpha}(\mathcal{S}) = p_i(\mathcal{S}; \pi_{i_\alpha})$*. Then we obtain the following necessary and sufficient condition for* $\mathcal{S}$ *to be an equilibrium point:*

$$p_i(\mathcal{S}) = \max_\alpha p_{i_\alpha}(\mathcal{S})$$

We can note that the above definition requires that $c_{i_\alpha} = 0$ whenever $p_{i_\alpha}(\mathcal{S}) < \max_\beta p_{i_\beta}(\mathcal{S})$, which means that $\mathcal{S}$ does not use π_{i_α} unless it is an optimal pure strategy for player i. Thus,

Definition 5.6 *If* π_{i_α} *is used in* $\mathcal{S}$ *then*

$$p_{i_\alpha}(\mathcal{S}) = \max_\beta p_{i_\beta}(\mathcal{S})$$

is another sufficient and necessary condition for an equilibrium point.

Again, Definitions 5.3, 5.5, and 5.6 all describe an equilibrium point.

Theorem 12 (The Fundamental Theorem) *Every finite game has an equilibrium point.*

Proof 12 We consider our definitions as above, so that $p_{i_\alpha}(\mathcal{S})$ is the payoff to player i if he changes to his α^{th} pure strategy $\pi_{\mathbf{i}_\alpha}$ and the others continue to use the mixed strategies from $\mathcal{S}$. We define a set of continuous functions of $\mathcal{S}$ by

$$\varphi_{i_\alpha}(\mathcal{S}) = \max(0, p_{i_\alpha}(\mathcal{S}) - p_i(\mathcal{S}))$$

we consider the mapping $T : \mathcal{S} \longrightarrow \mathcal{S}'$ which modifies each component of $\mathcal{S}$, $\mathbf{s_i}$, so that

$$\mathbf{s_i} \longmapsto \mathbf{s_i'} = \frac{\mathbf{s_i} + \sum_\alpha \varphi_{i_\alpha}(\mathcal{S})\pi_{i_\alpha}}{1 + \sum_\alpha \varphi_{i_\alpha}(\mathcal{S})}$$

So we call $\mathcal{S}'$ the n-tuple $(\mathbf{s_1'}, \mathbf{s_2'}, ..., \mathbf{s_n'})$.

We must now show that the fixed points of the mapping $T : \mathcal{S} \longrightarrow \mathcal{S}'$ are the equilibrium points.

We consider any n-tuple $\mathcal{S}$. In $\mathcal{S}$ the i^{th} player's mixed strategy $\mathbf{s_i}$ will use certain pure strategies. Some of these strategies must be least profitable, say π_{i_α} is such, then $p_{i_\alpha}(\mathcal{S}) \leq p_i(\mathcal{S})$. This will make $\varphi_{i_\alpha}(\mathcal{S}) = 0$.

If, however, the n-tuple $\mathcal{S}$ is fixed under T, the proportion of π_{i_α} used in $\mathbf{s}_i$ must not be decreased under T. Thus, $\forall$ β's, $\varphi_{i_\beta}(\mathcal{S})$ must be *zero* to prevent the denominator in our modification from exceeding 1.

We see that if $\mathcal{S}$ is fixed under T, for any i and for any β, $\varphi_{i_\beta}(\mathcal{S}) = 0$. Then, no player can improve his payoff by moving to a pure strategy π_{i_β}. By definition 5.5, this is exactly the criterion for an equilibrium point.

It is clear that if, conversely, $\mathcal{S}$ is an equilibrium point, then all the φ's become 0, making $\mathcal{S}$ a fixed point under T.

Finally, we can note that the n-tuple $\mathcal{S}$ lies in the space which is a product of the simplices with vertices π_{i_α} (each player has his own simplex). This space has dimension $d_1 + d_2 + ... + d_n$ where d_i represents the dimension of player i's simplex. But this space is homeomorphic to a ball of dimension $d_1 + d_2 + ... + d_n$, where it is clear that Brouwer's fixed point theorem applies. Hence, T has at least one fixed point $\mathcal{S}$, which we have shown to be an equilibrium point.

References

[1] Prisoner's Dilemma. Available at http://en.wikipedia.org/wiki/Prisoner

[2] Kuhn, Harold W. *Lectures on the Theory of Games.* Annals of Mathematics Studies. Princeton: Princeton University Press, 2003.

[3] Starr, Ross M. *General Equilibrium Theory.* Cambridge: Cambridge University Press, 1997.

[4] Taylor, Larry. "Sperner's Lemma, Brouwer's Fixed Point Theorem." Available at http://www.cs.csubak.edu/ larry/sperner.pdf.

[5] "The Mathematics of Strategy." Available at http://www.pnas.org/misc/classics5.shtml

[6] Nash, John. "Non-Cooperative Games." Annals of Mathematics. Vol. 54, No. 2 (September, 1951) pp. 286-295 .

[7] Tardos, Eva. "Randomized Nash Equilibrium." Available at http://www.cs.cornell.edu/courses/cs684/2005fa/10-26.pdf