AN ELEMENTARY EXPOSITION TO WILSON'S ODDNESS THEOREM

HRIDANK GARODIA

Abstract. The purpose of this paper is the provide an exposition of the Wilson's Oddness Theorem and make it easier for a non-expert to understand. This paper also provides a few exceptions to the theorem. It discusses a background in game theory, the history of game theory, the theorem itself, and Professor John C. Harsanyi's (University of California, Berkeley) proof to the theorem.

1. Introduction

Robert Wilson [Wil71] first displayed that apart from certain rare cases, in any finite game, the number of equilibrium cases are *finite* and *odd*. In simple, this is the Wilson's Oddness Theorem.

However, in this expository paper, we will be looking at John C. Harasanyi's proof to Wilson's Oddness Theorem [Har73]. This proof will be highlighted in sections 3, 4, 5, 6, 7, 8, and 9.

Let's look at some history of game theory in general and Wilson's Oddness Theorem in particular.

Discussions by Mathematicians on games began much before the modern Mathematical game theory. Cardano's Liber de ludo aleae (Book on Games of Chance), written about 1564 but published posthumously in 1663, articulated some of the field's fundamental notions. Several papers were published in this unofficial field until 1928.

In 1928, John von Neumann published his paper *On the Theory of Games of Strategy*. Von Neumann's original proof used Brouwer's fixed-point theorem on continuous mappings into compact convex sets, which became a standard method in game theory and mathematical economics.

This 1944 *Theory of Games and Economic Behavior* was the culmination of Von Neumann's work in game theory. The strategy for discovering mutually consistent solutions for two-person zero-sum games is described in this seminal paper. Following work concentrated mostly on cooperative game theory, which analyses optimum tactics for groups of persons under the assumption that they may enforce agreements among themselves concerning correct methods.

In 1950, the first mathematical discussion of the prisoner's dilemma appeared, and an experiment was undertaken by notable mathematicians Merrill M. Flood and Melvin Dresher, as part of the RAND Corporation's investigations into game theory.

Around the same time, John Nash devised the Nash equilibrium, a criteria for reciprocal consistency of players' tactics that is applicable to a broader range of games than the von Neumann and Morgenstern criterion. Nash demonstrated that every finite n-player, non-zero-sum (rather than only two-player zero-sum) non-cooperative game has a Nash equilibrium with mixed strategies.

Date: July 14, 2023.

Reinhard Selten proposed his solution notion of subgame perfect equilibria in 1965, which modified the Nash equilibrium further. He would later incorporate trembling hand perfection as well.

In 1971, Robert Wilson proved ***Wilson's Oddness Theorem*** about the number of nash equilibria. In 1973, John C. Harasanyi devised a new proof to Wilson's Oddness Theorem, which we are looking at in this paper. In 1994, Nash, Selten and Harasanyi became nobel laureates in Economics.

Next came evolutionry game theory by Thomas Schelling and Robert Aumann who were both awarded nobel prizes for their contributions.

This is a brief history of the evolution of game theory.

Acknowledgements

The author of this paper would like to thank Dr. Simon Rubenstein-Salzado and Mr. Alphonse Edige Houndegla for their contributions, such as conversation, to this paper.

2. Background

We will start off by providing a background game theory, in order to give context to inexperienced readers. If you feel you are confident in game theory, you may want to skip this section.

Game theory is the study of mathematical models of strategic interactions between rational agents . It has uses in computer science, logic, systems science, and all branches of social science. Economic theory also makes considerable use of game theory principles. In simpler words, game theory is the study and analysis of game situations.

We will now define some important terms in game theory and other terms for this paper. An article on Non-Cooperative Games [Nas51] by John F. Nash presented by Andrew Hutchings is definitely something the reader should consider going through for more depth. It will allow the reader to understand the following paper with more ease. It also provides some other theorems and proofs which the author finds quite interesting.

Definition 2.1. A *finite* $n-person$ game is a set of n players, each with an associated finite set of pure strategies.

Definition 2.2. A utility function U_i maps from the set of all n-tuples of pure strategies to the reals.

The utility function essentially gives the payoff of the strategy of the player in the game situation.

Definition 2.3. A mixed strategy of player i is a convex combination of pure strategies. Mixed strategies have a simple geometric representation as points on a simplex.

A mixed strategy is the probability distribution one uses to randomly choose among available actions in order to avoid being predictable and thus achieve the maximum payoff.

Additional explanation: Let $x_1, x_2, \cdots, x_n, x_{n+1}$ be $n+1$ points in R^n. Then the convex set spanned by these $(n+1)$ points is called an n-simplex in R^n. For example, a 0-simplex is a point, a 1-simplex is a line, a 2-simplex can be a triangle, and a 3-simplex can be a tetrahedron. If you do not understand this information, I recommend checking out a book on algaebric topology since the proof will require some of it.

Note: The aim of game theory is usually to analyse a situation in order to understand the best possible strategy to maximize payoff or U_i.

Definition 2.4. Equilibrium Point: A $n-$tuple S is an equilibrium point if and only if for every i:

$$U_i(s) = \underset{allr_i's}{max}[U_i(S; r_i)]$$

which by linearity of p_i gives us:

$$\underset{\alpha}{max}[U_i(S; \pi_{i\alpha})] = \underset{allr_i's}{max}[U_i(S; r_i)]$$

.

Define $p_{i\alpha}(S) = p_i(S; \pi_{i\alpha})$.

Then we have S is an equilibrium point if:

$$U_i(S) = \underset{\alpha}{max} U_{i\alpha}(S)$$

.

In Definition 2.4, i indicates a player, α indicates a pure strategy of a player, r_i indicates a mixed strategy of i and π_{ia} indicates i's α^{th} pure strategy.

3. Defining Key Terms

First we look at pure strategies,

The $k-th$ pure strategy of player $i(i = 1, \dots, n)$ will be called a_i^k and the set of all the player's K_i pure strategies will be called A_i. Let

$$K = \prod_{i=1}^{n} K_i \tag{3.1}$$

We shall assume that the K possible $n-tuples$ of pure strategies are numbered consecutively as $a^1, \dots, a^m, \dots, a^k$. Let

$$a^m = (a_1^{k_1}, \dots, a_i^{k_i}, \dots, a_n^{k_n}). \tag{3.2}$$

Thus it follows that

$$a^m(i) = (a_i^{k_i}), \tag{3.3}$$

denotes the pure strategy used by player i in the strategy n-tuple a^m. The set of all K pure strategy n-tuples will be called A. Thus, $A = A_1 * \cdots * A_n$.

Now we look at mixed strategies,

Any mixed strategy n-tuple of a given player $i(1,\ldots,n)$ can be identified with a probability vector p_i as

$$p_i = (p_i^1, \ldots, p_i^k, \ldots, p_i^{Ki}) \tag{3.4}$$

wherein p_i^1 is the probability that this mixed strategy assigns to the player's pure strategy a_i^1. (Definition 2.3 of a mixed strategy is relevant in this case).

As mentioned in Definition 2.3, since mixed strategies can be represented as points on a simplex; The set P_i of all mixed strategies available to i is a simplex consisting of all K_i-vectors, satisfying the conditions:

$$p_i^k \geq 0 \text{ for } k = 1, \ldots, K_i, \tag{3.5}$$

and

$$\sum_{k=1}^{K_i} p_i^k = 1. \tag{3.6}$$

The set $P = P_1 * \cdots * P_n$ of all n-tuples $p = (p_i, \cdots, p_n)$ of mixed strategies forms a compact and complex polyhedron. We will call this the *strategy space* of game Γ. We shall write $p = (p_i, \overline{p_i})$, wherein we write $\overline{p_i} = (p_i, \ldots, p_{i-1}, p_{i+1}, \ldots, p_n)$ is the strategy $(n-1)$-tuple representing the mixed strategies of $(n-1)$ players other than the player i.

Assigning Pure Strategies to Mixed Strategies,

For any strategy n-tuple $p = (p_1, \ldots, p_n)$ the carrier of $C(p)$ will be defined as the union of carrier of its component strategies as:

$$C(p) = \bigcup_{i=1}^{n} C(p_i). \tag{3.7}$$

Additional Explanation: The carrier of p_i is the set $C(p_i)$ of all pure strategies a_i^k to which the mixed strategy p_i assigns positive probabilities $p_i^k > 0$. If this carrier contains only one pure strategy a_i^k, then $a_i^k = p_i$. However, if $C(p_i)$ contains all K_i pure strategies of player i, then p_i will be called a *complete mixed* strategy. And lastly, if p_i doesn't have only one or all pure strategies, it will be called an *incomplete mixed* strategy.

Now we will assume that the i^{th} component of the pure strategy n-tuple a^m is $a^m(i) = a_i^k$ and that a given mixed strategy p_i of player i assigns the mixed strategy p_i^k to the pure strategy a_i^k.

From this we can write

$$q_i^m(p_i) = p_i^k. \tag{3.8}$$

Now, with the help of equations 3.7 and 3.8, if $p_i = a_i^k$ is a pure strategy, then

$$q_i^m(a_i^k) = 1 \text{ when } a^m(i) = a_i^k, \tag{3.9}$$

$$q_i^m(a_i^k) = 0 \text{ when } a^m(i) \neq a_i^k. \tag{3.10}$$

This is again quite logical, since if it is a pure and only strategy, the probability of assigning to it will be 1. And the probability of assigning to it when $a^m(i) \neq a_i^k$ will be 0 because $a^m(i)$ won't be equal to the i^{th} component of the pure strategy. In simpler terms, we will be trying to assign to something that doesn't exist.

When the n players use pure strategy n-tuple a^m, then player $i(i = 1, 2, \ldots, n-1, n)$ will receive the payoff (as per the utility function in Definition 2.2)

$$U_i(a^m) = u_i^m, \tag{3.11}$$

whereas it's a little more complicated for the using the mixed strategy n-tuple $p(p_1, \ldots, p_n)$. His payoff with it will be

$$U_i(p) = \sum_{m=1}^{K} [\Pi_{i=1}^n q_i^m(p_i)] u_i^m \tag{3.12}$$

which utilises the payoff for the pure strategy

Let $\zeta = \zeta(n; K_1, \ldots, K_n)$ be the set of all n-person games in which players $1, \ldots, n$ have exactly $K_1, \ldots, K_n$ pure strategies, respectively. Thus, ζ is the set of *all* games of a given size. Each specific game Γ in the set of all games ζ can be characterised by the (nK)-vector

$$u = (u_1^1, \ldots, u_1^{K_1}; \ldots; u_i^1; \ldots, u_i^{K_i}; \ldots; u_n^1; \ldots, u_n^{K_n}) \tag{3.13}$$

whose components $u_i^m = U_i(a^m)$ are the payoff to various players i for the different pure strategies combinations a^m. We can identify each game Γ with its vector $u = u(\Gamma)$ of possible payoffs of pure-strategy combinations, and can regard the set of all games ζ as an (nK)-dimensional Euclidean space $\zeta = u$.

Important Preliminary to Wilson's Oddness Theorem Proof

Let $\overline{x}(\zeta)$ be the set of all games Γ in ζ for which a given mathematical statement x is false. This proof for Wilson's Oddness Theorem relies on saying that statement x is true for *almost all* games, if, for every possible set ζ of games of a particular size, this set $\overline{x}(\zeta)$ is a *closed* set of *measure zero* within the relevant set ζ, regarded as an nK - dimensional Euclidean Space. The closure requirement is cited from [Deb70]

4. Logarithmic games

Harasanyi's paper [Har73] uses an innovative method of depicting logarithmic games in the following way.

Let Λ be an n-person 'non cooperative game', where the n players have the same simplex $P_1, \ldots, P_n$ they have in game Γ as strategy spaces, but where the payoff function L_i of each player $i(i = 1, \ldots, n)$ is of the form

$$L_i(p) = L_i(p_i) = \sum_{k=1}^{K_i} \log p_i \tag{4.1}$$

Since the payoff function L_i are logarithmic functions , this game Γ - as well as the games $\Gamma * (t)$ to be defined below - are best regarded as being *infinite* games in which the *pure* strategy of every player i consists in choosing a specific point p_i from the simplex P_i, which makes p_i a pure strategy and not mixed. But, for convenience, Harasanyi goes on calling any given strategy p_i a (complete or incomplete) mixed strategy. This terminology will not give rise to any confusion because, in analyzing these infinite games, we shall never consider mixed strategies having the nature of probability mixtures of two or more p_i.

Thus, Λ is a "degenerate".

Note: A "degenerate" game is a game wherein each players payoff L_i depends only on his own strategy p_i and not on other player's strategies p_j, for $j \neq i$

We define a one-parameter family of games $A * (t)$, with $0 \leq t \leq 1$. In any particular game $\Lambda * (t)$ with a specific value of the parameter t, the payoff function of player $i (i = 1, \ldots, n)$ is

$$L * (p, t) = (1 - t)U_i(p) + tL_i(p_i). \tag{4.2}$$

$\Lambda * (0) = \Gamma$, whereas $\Lambda * (1) = \Lambda$. All games $\Lambda * (t)$ with $0 < t \leq 1$ will be called *logarithmic games*. Γ will be called the *original game* and Λ will be called the *pure logarithmic* game.

What we have shown in this section, is how Harasanyi classifies logarithmic games very intelligently.

5. Conditions

There exists a best reply strategy p_i of a player i to a strategy combination $\overline{p_i}$, used by the other $(n - 1)$ players in the game $A * (t)$ if

$$L * (p_i, \overline{p_i}, t) \geq L * (p'_i, \overline{p_i}, t) \text{ for all } p'_i \in P_i. \tag{5.1}$$

Referring to Definition 2.4 [Nas51], we can also classify an equilibrium point if every component p_i of p is the best corresponding reply to a strategy by the $(n - 1)$ players.

Definition 5.1. Strong equilibrium point: If all n components of p_i of p satisfy with the strong equality sign $>$ for all $p'_i \neq p_i$.

Note: p is a strong equilibrium point if every player's equilibrium strategy p_i is his *only* best reply to the other players' strategy combination $\overline{p_i}$.

Definition 5.2. Weak equilibrium point: If the equilibrium point is not strong, then it is weak.

Definition 5.3. Quasi Strong Equilibrium Point: p is quasi strong if no player i has has pure-strategy best replies to $\overline{p_i}$ other than the pure strategies belonging to carrier $C(p_i)$ of his equilibrium point p_i.

Definition 5.4. Extra Weak Equilibrium Point: An equilibrium point that is not even quasi-strong is called extra-weak.

We can classify games such that Γ will be a quasi strong game if all equilibrium points are quasi strong. However, it will be extra weak if one or more are extra weak.

In the original game $\Gamma = A*(0)$, a best reply p_i to any given strategy combination $\overline{p_i}$ may be a pure strategy or may be a mixed strategy. (It can be a mixed strategy only if all pure strategies a_i^k in its carrier $C(p_i)$ are themselves best replies to p_i.)

However, if we look at the logarithmic game $\Lambda*(t)$ with $t > 0$, only a *complete* mixed strategy can be the best reply. This is because from equations (3.12), (4.1) and (4.2) any player i will obtain a negative payoff if they use a pure strategy or an incomplete mixed strategy.

In the original game Γ, in general, the mathematical conditions characterizing an equilibrium point $p = (p_1, \ldots, p_n)$ will be partly equations and partly inequalities. The equations of the equilibrium point of Γ will be of the form:

$$U_i(a_i^k, \overline{p_i}) = U_i(a_i'^k, \overline{p_i}) \text{ if } a_i^k, a_i'^k \in C(p_i); \tag{5.2}$$

whereas the inequalities will be:

$$U_i(a_i^k, \overline{p_i}) \geq U_i(a_i'^k, \overline{p_i}) \text{ if } a_i^k \in C(p_i), \text{ while } a_i'^k \notin C(p_i). \tag{5.3}$$

The case where all n equilibrium strategies are pure strategies will the inequalities (5.3) will be characterised. The case where all all n equilibrium strategies are complete mixed strategies will characterise the equation (5.2).

An equilibrium point p in the finite game Γ will be quasi strong if and only if, for every player i, and for every strategy a_i^k in $C(p_i)$ and for every strategy $a_i^{k'}$ not in this carrier, inequality (5.3) is followed with a strong equality sign.

In contrast, every equilibrium point p in the logarithmic $A*(t)$ with $t > 0$ is always characterised by these equations.

$$\frac{\partial L*_i(p,t)}{\partial p_i^k}_{\sum_i p_i^k = 1} = 0, \text{ for } , k = 1, \ldots, K_i - 1, \text{ and, for, } i = 1, \ldots, n. \tag{5.4}$$

Additional Explanation: The partial derivative has been taken to maximise the payoff function $L*_i$ with respect to vector p_i

Each partial derivative in (5.4) must be equated at the equilibrium point p itself. Since each function $L*_i$ is strictly concave in p_i, the second-order conditions are always satisfied, so that the eqs. (5.4) are both necessary and sufficient conditions for maximization. The function $L*_i$ can also be written as:

$$L*_i(p,t) = (1-t)\sum_k^{K_i} = 1 p_i^k U_i(a_i^k, \overline{p_i}) + \sum_k^{K_i} = 1 log p_i^k. \tag{5.5}$$

Using the fact:

$$p_i^1 = 1 - \sum_k^{K_i} = 2p_i^k, \tag{5.6}$$

We can write (5.4) in the form:

$$(1-t)[U_i(a_i^k, \overline{p_i}) - U_i(a_i^1, \overline{p_i})] + \frac{t}{p_i^k} - \frac{t}{p_i^1} = 0, \tag{5.7}$$

or, equivalently in the form

$$(1-t)p_i^1 p_i^k [U_i(a_i^k, \overline{p_i}) - U_i, (a_i^1, \overline{p_i})] + t(p_i^1 - p_i^k) = 0, \text{for, } i = 1, \ldots, n; \text{and, for, } k = 2, \ldots, K_i. \tag{5.8}$$

Additional Explanation: Going from equation 5.7 to equation 5.8 essentially requires taking t out common and multiplying the whole equation by $p_i^1 p_i^k$.

The number of equations of form (5.8) can be written as:

$$K* = \sum_{i=1}^{n} (K_i - 1) = \sum_{i=1}^{n} K_i - n. \tag{5.9}$$

Looking at n equations of the form (3.6) we have all together:

$$K** = K* + n = \sum_{i=1}^{n} K_i \tag{5.10}$$

number of independent equations for characterising each equilibrium point p, which is the same number of the variables p_i^k determined by these equations.

6. Algebraic Curves

Definition 6.1. An algebraic curve over a field K is an equation $f(X, Y) = 0$, where $f(X, Y)$ is a polynomial in X and Y with coefficients in K. A nonsingular algebraic curve is an algebraic curve over K which has no singular points over K. A point on an algebraic curve is simply a solution of the equation of the curve. A K-rational point is a point (X, Y) on the curve, where X and Y are in the field K.

In view of (3.12), all equations of the form (5.8) are algebraic equations in the variables p_i^k and parameter t.

All equations of form (3.6) are likewise algebraic. Let S be the set of all $K** + 1$-vectors (t, p) satisfying the $K**$ equations of form (3.6) and (5.8).

Note: S will be typically a one-dimensional algebraic variety, i.e., an algebraic curve. (In degenerate cases, however, S may also contain zero-dimensional subsets, i.e., isolated points, and/or subsets of more than one dimension, i.e., algebraic surfaces of various dimensionality.)

Now, we define T as the set of all vectors (t, p) that satisfy not only the $K**$ equations of form (3.6) and (5.8) but also the $K**$ inequalities of form (3.5).

Note: T is simply that part of the algebraic variety S which lies within the compact and convex cylinder (polyhedron) $R = P * I$, where $I = [0, 1]$ is the closed unit interval.

Since T is the locus of all solutions (t, p) to the simultaneous equations and inequalities (3.5), (3.6), and (5.8), T will be called the solution graph for the latter.

Important for upcoming sections: For any point (t, p), t will be called its first coordinate. Within the cylinder R, the strategy space P of any specific game $A * (t)$ is represented by the set R^t of all points (t, p) in R.

For any game $A * (t)$, let E^t be the set of all points (t, p) in R^t such that p is an equilibrium point of $A * (t)$. Finally, let T^t be the intersection of R^t with the solution graph T.

Lemma 6.2. *For all t with $0 < t \leq 1, E^t = T^t$. In contrast, for $t = 0$, in general, we have only $E^0 \subseteq T^0$.*

Proof. For all t with $0 < t \leq 1$, conditions (3.5), (3.6), and (5.8) are sufficient and necessary conditions for any given point p to be an equilibrium point of game $A * (t)$. However, for $t = 0$, it is easy to see that all equilibrium points p of the game $A * (0) = \Gamma$ satisfy all these conditions but, in general, so will strategy combinations that are not equilibrium points of Γ. For example, all these conditions will be satisfied by any pure-strategy n-tuple $p = a^m$, whether it is an equilibrium point of game Γ or not.

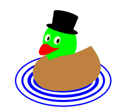

7. Topological Properties of Solution Graph T

Consider the mapping $\mu : t \to T^t$. The Jacobian of this mapping, as evaluated at any given point (p, t) of T^t, can be written as:

(7.1)
$$J(t, p) = \frac{\partial(F_1^1, \ldots, F_i^k, \ldots, F_n^{K_n})}{\partial(p_1^1, \ldots, p_i^k, \ldots, p_n^{K_n})}$$
for $i = 1, \ldots, n$;
and for each $i, k = 1, \ldots, K_i$.

Here:

(7.2)
$$F_i^1(t, p) = \sum_{i=1}^{n} p_i^k - 1, \qquad i = 1, \ldots, n;$$

whereas,

(7.3)
$$F_i^k(t, p) = (1 - t) p_i^1 p_i^k [U_i(a_i^k, \overline{p_i}) - U_i(a_i^1, \overline{p_i})] + t(p_i^1 - p_i^k)$$
$i = 1, \ldots, n$;
and for each i, $\quad k = 2, \ldots, K_i$.

For points of the form $(t, p) = (0, p)$ in set T^0 the functions $F_i^k (k \neq 1)$, the following simpler form of (7.3) arises:

(7.4)
$$F_i^k(0, p) = p_i^1 p_i^k [U_i(a_i^k, \overline{p_i})],$$
for $i = 1, \ldots, n$;
and, for each i, $k = 2, \ldots, K_i$.

Definition 7.1. An equilibrium point of game Γ will be called regular if $J(0,p) \neq 0$; and will be called irregular if $J(0,p) = 0$.

Definition 7.2. A given game Γ will be regular if all of its equilibrium points are regular, and will be irregular if *one* or more of its equilibrium points are irregular.

Now we will state two lemmas based on well known facts in algebraic geometry.

Lemma 7.3. *Let (x^0, x^*) be an arc of an algebraic curve S in a v-dimensional Euclidean space X^v, connecting two points $x^0 = (x_1^0, \dots, x_v^0)$ and $x^* = x_1^*, \dots, x_v^* \neq x^0$. Then, this arc (x^0, x^*) can be uniquely continued analytically beyond point x^* (and beyond point x^0).*

Proof. If x^* is not a singular point, then the possibility of analytic continuation follows from the Implicit Function Theorem. On the other hand, if x^* is a singular point, then this possibility follows from PUISEUX's Theorem [Kun97] Theorem 14. By this theorem, if x^* is a point of some branch S^* of a given algebraic curve S, then, whether x^* is a singular point or not, in some neighborhood $N(x^*)$ of x^*, the coordinates x_i of any point $x = (x_1, \dots, x_v)$ of this branch $S*$ can be represented by v convergent power series $\pi_i(y)$ in an auxiliary parameter y, so that we can write $x_i = \pi_i(y)$ for a suitably chosen value of $y (i = 1, \dots, v)$. Moreover, we can select the v functions n, in such a way that x^* itself will correspond to $y = 0$ (so that $x_i^* = \pi_i(0)$, for $i = 1, \dots, v$), and in such a way that all other points x of the arc (x^0, x^*) will correspond to negative values of y. Then, by assigning positive values to y, we can analytically continue the arc (x^0, x^*) beyond x^*. Even though we can choose the v functions π_i in many different ways, all choices will yield the same curve as the analytic continuation of (x^0, x^*).

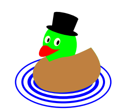

Corollary 7.4. *Let S be an algebraic curve, and x be an arbitrary point. Then, the number of arcs belonging to S and originating from x is always even (possibly zero). These arcs always uniquely partition themselves into pairs, so that the two arcs belonging to the same pair are analytic continuations of each other, and are not analytic continuations of any other arc originating from x.*

Lemma 7.5. *Let (x^0, x^*) be an arc of an algebraic curve S. Suppose that $(x^0 x^*)$ lies wholly within a given compact and convex set R with a nonempty interior, and that x^0 is a boundary point of R whereas x^* is an interior point of R. Then, by analytically continuing (x^0, x^*) far enough beyond x^*, we shall once more eventually reach a boundary point x^{00} of R.*

Proof. We define S^* as the algebraic curve if we extend x^* to a boundary point x^{00} of R. For each coordinate x_i, let

$$m_i = \underset{x \in S^*}{inf} x_i \quad and \quad m^i = \underset{x \in S^*}{sup} x_i. \tag{7.5}$$

Since S^* is not an isolated point, at least for one coordinate x_i, its variation on $S^*, \Delta_i = m^i - m_i$, must be positive. On the other hand, since S^* is an arc of an algebraic curve, it can be divided up into a finite number of segments $\alpha^1, \dots, \alpha^\mu, \dots, \alpha^M$, such that, as we move away from x^0 along any given segment α^μ, this coordinate x_i is either strictly increasing or is strictly decreasing. Let us assume that, starting from x^0 and moving along S^*, we reach these segments in the order they have been listed. Now, first suppose that, along the last segment α^M, x_i increases. Then, since R is a compact set, x_i must reach a local maximum at some point x^{00} of α^M. Obviously, this point x^{00} can only be the endpoint of α^M furthest away from x^0. Moreover, it can only be a boundary maximum point for x_i because, if it were an interior maximum point, then α^M could not be the last segment of S^*. Therefore, this point x^{00} must be a boundary point of R. By the same token, if x_i decreases along α^M, then the endpoint x^{00} of α^M must be a local boundary minimum point for x_i and, therefore, it must be boundary point of R. Thus, in either case, S^* will eventually reach a boundary point x^{00} of R.

Let $\bar{P}$ be the boundary of the strategy space P. Thus, $\bar{P}$ is the set of all strategy n-tuples $p = (p_1, \ldots, p_n)$ having at least one pure or incompletely mixed strategy p_i as a component. Let $I^0 = (0, 1)$ be the open unit interval. Let B be the set $B = \bar{P} * I^0$. Clearly, the boundary hypersurface $\bar{R}$ of cylinder R is made up of the three disjoint sets P^0, B, and P^1.

Lemma 7.6. *Let $\bar{T}$ be the intersection of the solution graph T and of the boundary hypersurface $\bar{R}$ of cylinder R. Let (t, p) be any nonisolated point of $\bar{T}$. Then, either $(t, p) = (1, \tilde{p})$, where $\tilde{p}$ is the unique equilibrium point of the pure logarithmic game $\Lambda^*(1)$; or $(t, p) = (0, p^*)$, where p^* is an equilibrium point of the original game $\Lambda^*(0) = \Gamma$*

Proof. For any t with $0 < t \leqq 1$, the vector p characterizing any given point (t, p) of $\bar{T}$ must be an equilibrium point of the game $\Lambda^*(t)$, because (t, p) is a point of the solution graph T. Therefore, (t, p) cannot belong to set B, since the logarithmic games $\Lambda^*(t)$ with $0 < t \leqq 1$ have no equilibrium point using a pure or an incompletely mixed strategy p_i as equilibrium strategy. Hence, if $t > 0$, then (t, p) can only be a point belonging to set P^1, which is possible only if $(t, p) = (1, \tilde{p})$. On the other hand, if $t = 0$, then $(t, p) = (0, p)$ is a point belonging to set P^0. As (t, p) is a nonisolated point of T, it is a limit point of some convergent point sequence $(t^1, p^1), \ldots, (t^j, p^j), \ldots$, where each p^j is an equilibrium point in game $\Lambda^*(t^j)$, with $t^j > 0$. Consequently, p itself is an equilibrium point in game $\Lambda^*(0) = \Gamma$, because the correspondence $\mu^* : t \to E^t$ is upper semi-continuous (where E^t is the set mentioned in Lemma 1). This completes the proof.

Lemma 7.7. *The point $(t, p) = (1, p)$, corresponding to the unique equilibrium point $\tilde{p}$ of the pure logarithmic game $\Lambda^*(1) = \Lambda$ is always a nonsingular point of the graph T, and is the endpoint of exactly one branch $\beta(\tilde{p})$ of T.*

Proof. As is easy to verify, $J(1, \tilde{p}) \neq 0$. Consequently, $(1, \tilde{p})$ is nonsingular and, by the Implicit Function Theorem, it lies on exactly one branch $\alpha(\tilde{p})$ of T.

Lemma 7.8. *Let Γ be a regular and quasi-strong game. Then, any point $(t, p) = (0, p)$ corresponding to an equilibrium point p of game $\Lambda^*(0) = \Gamma$ is always a nonsingular point of the graph T, and is the endpoint of exactly one branch $\beta(p)$ of T.*

Proof. Since Γ is regular, we have $J(0, p) \neq 0$. Hence, if p is an interior point of the strategy space P, then the present lemma can be established by the same reasoning as was used in the proof of Lemma 5. However, if p is a boundary point of P, then this reasoning shows only that $(0, p)$ lies on exactly one branch $\beta(p)$ of the algebraic variety S. In order to prove the lemma, we have to show also that $\beta(p)$ belongs to the graph T, i.e., that it lies within cylinder R. In other words, we have to show that $\beta(p)$ goes from $(0, p)$ towards the interior of R, which is equivalent to showing that, for any zero component $p_i^k = 0$ of the vector p, the total derivative $\mathrm{d}p_i^k/\mathrm{d}t$ is positive at the point $(0, p)$. Now, by differentiating eq. (5.8) with respect to t, and then setting $t = p_i^k = 0$, we obtain

$$p_i^1 \left[U_i\left(a_i^k, \bar{p}_i\right) - U_i\left(a_i^1, \bar{p}_i\right)\right] + p_i^1 = 0. \tag{7.6}$$

Since the numbering of player i 's pure strategies is arbitrary, without loss of generality we can assume that

$$p_i^1 > 0 \tag{7.7}$$

On the other hand, since $p_i^1 > 0$ and $p_i^k = 0$, we have $p_i^1 \in C\left(p_i\right)$ but $p_i^k \notin C\left(p_i\right)$. Since p is a quasi-strong equilibrium point, condition (18) must be satisfied by a strong inequality sign if we set $a_i^{k'} = a_i^1$. Therefore,

$$U_i\left(a_i^k, \bar{p}_i\right) - U_i\left(a_i^1, \bar{p}_i\right) < 0. \tag{7.8}$$

But (7.6), (7.7), and (7.8) together imply that $\mathrm{d}p_i^k/\mathrm{d}t > 0$, as desired.

In what follows, when we say that two points are "connected", we shall mean that they are connected by some branch α of the solution graph T.

Theorem 7.9. *Let Γ be a regular and quasi-strong finite game. Then, the number of equilibrium points in Γ is finite. Moreover, there exists exactly one distinguished equilibrium point p^* in Γ such that the corresponding point $(0, p^*)$ is connected with the point $(1, \tilde{p})$, associated with the unique equilibrium point $\tilde{p}$ of the pure logarithmic game $\Lambda^*(1) = \Lambda$. All other equilibrium points of Γ form pairs, such that the two equilibrium points belonging to the same pair are connected with each other and with no other equilibrium point. Therefore, the number of equilibrium points in Γ is odd.*

Proof. By Lemma 6, every equilibrium point p of Γ lies on some branch $\beta(p)$ of T. But T, being the intersection of an algebraic variety S and of a compact and convex set R, can have only a finite number of branches. Moreover, on any given branch β, there can be at most two equilibrium points, corresponding to the two endpoints of β. Therefore, the number of equilibrium points in Γ is finite.

By Lemma 5 , there exists a unique branch $\alpha(p)$ of T, originating from the point $(1,\tilde{p})$. By Lemmas 1 and 2, this branch $\alpha(\tilde{p})$ must lead to a boundary point x^{00} of R. As $J(1,\tilde{p}) \neq 0$, we must have $x^{00} \neq (1,\tilde{p})$, because otherwise T would have two local branches originating from $(1,\tilde{p})$, contrary to the Implicit Function Theorem. Consequently, by Lemma 4, $x^{00} = (0,p^*)$, where p^* is an equilibrium point - called the distinguished equilibrium point - of game Γ.

Finally, let $p \neq p^*$ be any equilibrium point of Γ, other than the distinguished equilibrium point p^*. By Lemma 6 , there exists a unique branch $\beta(p)$ of T, originating from the point $(0,p)$. By an argument similar to the one used in the last paragraph, it can be shown that $\beta(p)$ must lead to another boundary point x^{00} of R, with $x^{00} = (0,p')$, where $p' \neq p$ and $\neq p^*$ is another equilibrium point of Γ. Hence, all equilibrium points of Γ, other than the distinguished equilibrium point $\tilde{p}$, are pairwise connected. But this means that the number of these latter equilibrium points is even, which makes the total number of equilibrium points in Γ odd.

Note: The proof of Theorem (7.9) shows that, for any game $\Lambda^*(t)$ with $0 \leqq t \leqq 1$, the set Q^t of all equilibrium points in $\Lambda^*(t)$ is nonempty. This is so because branch $\alpha(\tilde{p})$ of graph T connects the two points $(1,\tilde{p})$ and $(0,p^*)$. Therefore, $\alpha(\tilde{p})$ intersects every set R^t with $0 \leqq t \leqq 1$ at some point (t,p^t). As is easy to verify, the strategy n-tuple p^t defining this point must be an equilibrium point of game $\Lambda^*(t)$.

8. Finite Games Being Quasi Strong

Within a given set $\mathscr{I} = \mathscr{I}(n; K_1, \ldots, K_n)$ of games of a particular size, let $\mathscr{F}(C^*)$ be the set of all games Γ that have at least one equilibrium point p with a specified set $C^* = C(p)$ as its carrier. There are only a finite number of different sets $\mathscr{F}(C^*)$ in $\mathscr{I}$ because, for all games Γ in $\mathscr{I}$, the number of possible carrier sets C^* is finite. This is so because any set C^* is a subset of the finite set

$$A^* = \bigcup_{i=1}^{n} A_i \tag{8.1}$$

consisting of the set of all K^{**} pure strategies a_i^k for the n players in each game Γ, where

$$K^{**} = \sum_{i=1}^{n} K_i \tag{8.2}$$

(Of course, two sets $\mathscr{F}(C^*)$ corresponding to different carrier sets C^* will in general overlap.) We can now state the following theorem.

Theorem 8.1. *Almost all finite games are quasi-strong.*

Proof. Let $\overline{\mathscr{F}}(C^*)$ be the set of all games Γ in $\mathscr{I}$ that have at least one extra-weak equilibrium point p with the set $C^* = C(p)$ as its carrier. Obviously, $\overline{\mathscr{F}}(C^*)\, C \mathscr{F}(C^*)$. Let $\overline{\mathscr{F}}(C^*) = \mathscr{F}(C^*) - \overline{\mathscr{F}}(C^*)$. Thus, all games Γ in $\overline{\mathscr{F}}(C^*)$ have the property that they contain one or more equilibrium points p with the set $C^* = C(p)$ as their carrier set, and all these equilibrium points p are quasi-strong.

All games Γ in a given set $\mathscr{F}(C^*)$ are characterized by the fact that their defining vector $u = u(\Gamma)$ satisfies a finite number of algebraic equations and algebraic weak inequalities, of forms (5.2) and (5.3), in which the functions U_i are defined by (3.12). Thus, if we regard the set $\mathscr{I}$ as an (nK)-dimensional Euclidean space $\mathscr{I} = \{u\}$, then each set $\mathscr{F}(C^*)$ will correspond to a subset of $\mathscr{I}$, bounded by pieces of a finite number of algebraic hypersurfaces. In view of (3.12) and (5.3), these bounding hypersurfaces are multilinear, i.e., they are hyperboloids. Within each set $\mathscr{F}(C^*)$, all games Γ belonging to $\overline{\mathscr{F}}(C^*)$ are characterized by the fact that their defining vectors $u = u(\Gamma)$ satisfy all the inequalities of form (5.3) used in defining this set $\mathscr{F}(C^*)$, with a strong inequality sign $>$. In contrast, all games Γ belonging to $\overline{\mathscr{F}}(C^*)$ have a defining vector $u = u(\Gamma)$ satisfying one or more of these weak inequalities with an equality sign $=$. Therefore, all games in $\overline{\mathscr{F}}(C^*)$ correspond to interior points u of $\mathscr{F}(C^*)$, whereas all games in $\overline{\mathscr{F}}(C^*)$ correspond to boundary points of $\mathscr{F}(C^*)$. Hence, as a subset of the (nK)-dimensional Euclidean space $\mathscr{I}, \overline{\mathscr{F}}(C^*)$ consists of pieces of a finite number of hyperboloids of at most $(nK-1)$ dimensions. Consequently, each set $\overline{\mathscr{F}}(C^*)$ is a set of measure zero in $\mathscr{I}$. Let $\overline{\mathscr{F}}*$ be the set of all extra-weak games in $\mathscr{I}.\overline{\mathscr{F}}$ is the union of all sets $\overline{\mathscr{F}}(C^*)$, corresponding to various possible carrier sets C^*. Thus, $\overline{\mathscr{F}}*$ is a union of a finite number of sets of measure zero in $\mathscr{I}$. Therefore, $\overline{\mathscr{F}*}$ itself is also a set of measure zero in $\mathscr{I}$.

Next, we shall show that $\overline{\mathscr{F}}^*$ is a closed set. Let $\Gamma^1, \Gamma^2, \ldots$ be a sequence of extra-weak games in $\mathscr{I}$, with the defining vectors $u^1 = u(\Gamma^1), u^2 = u(\Gamma^2), \ldots$. Suppose that the sequence $u^1, u^2, \ldots$ converges to a given vector u^0. Let Γ^0 be the game corresponding to this vector $u^0 = u(\Gamma^0)$. We have to show that Γ^0 is likewise an extra-weak game.

Since the games $\Gamma^j (j = 1, 2, \ldots)$ are extra-weak, each vector u^j satisfies one or more inequalities of form (5.3), with an equality sign. Yet, there are only a finite number of inequalities of this form. Therefore, at least one of these inequalities - let us call it inequality $(5.3)^*$ - will be satisfied by infinitely many vectors u^j, with an equality sign. As the sequence of these latter vectors, being a subsequence of the original sequence $\{u^j\}$, converges to u^0, this vector u^0 itself will also satisfy $(5.3)^*$ with an equality sign, which makes the corresponding game Γ^0 extra-weak, as desired. This completes the proof of Theorem 2 .

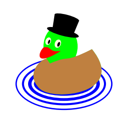

Let p be an equilibrium point in game Γ belonging to set $\mathscr{I}$, with the carrier $C^* = C(p) = \bigcup_i C(p_i)$. Thus $\Gamma \in \mathscr{F}(C^*)$. Suppose the carriers $C(p_1), \ldots, C(p_n)$ of the equilibrium strategies $p_1, \ldots, p_n$ consist of exactly $\gamma_1, \ldots, \gamma_n$ pure strategies, respectively. In studying games Γ in set $\mathscr{F}(C^*)$, we shall adopt the following notational convention, which, of course, involves no loss of generality:

(1) The pure strategies a_i^k of each player $i (i = 1, \ldots, n)$ have been renumbered in such a way that the carrier $C(p_i)$ of his equilibrium strategy p_i now contains his first γ_i pure strategies $a_i^1, \ldots, a_i^{\gamma_i}$.

We can fully characterize each equilibrium strategy p_i by the $(\gamma_i - 1)$ probability numbers $p_i^2, p_i^3, \ldots, p_i^{\gamma_i}$, since we have

$$p_i^1 = 1 - \sum_{j=2}^{\gamma_i} p_i^j \tag{8.3}$$

and

$$p_i^k = 0, \quad \text{for} \quad k = \gamma_i + 1, \ldots, K_i. \tag{8.4}$$

Let π_i be the probability vector

$$\pi_i = \left(p_i^2, \ldots, p_i^{\gamma_i}\right), \quad \text{for} \quad i = 1, \ldots, n. \tag{8.5}$$

Thus, π_i is a subvector of the probability vector p_i. Let $\Pi_i(i = 1, \ldots, n)$ be the set of all $(\gamma_i - 1)$-vectors satisfying the two conditions

$$p_i^k > 0, \quad k = 2, \ldots, \gamma_i \tag{8.6}$$

and

$$\sum_{j=2}^{y_i} < 1 \tag{8.7}$$

Let π be the composite vector

$$\pi = (\pi_1, \ldots, \pi_n). \tag{8.8}$$

Thus, π is a vector consisting of γ^* probability numbers p_i^k, where

$$\Gamma^* = \sum_{i=1}^{n} (\gamma_i - 1) = \sum_{i=1}^{n} \gamma_i - n \tag{8.9}$$

Clearly, π is a subvector of the probability vector p. Let Π be the set of all γ^*-vectors π whose subvectors $\pi_1, \ldots, \pi_n$ satisfy conditions (8.6) and (8.7). Clearly, $\Pi = \Pi_1 \times \cdots \times \Pi_n$. We now define

$$m^*(1, k) = k - 1, \quad \text{for} \quad k = 2, \ldots, \gamma_1 \tag{8.10}$$

and

$$\begin{aligned} m^*(i, k) &= \sum_{j=1}^{i-1} (\gamma_j - 1) + (k - 1) \\ &= \sum_{j=1}^{i-1} \gamma_j - i + k, \quad \text{for} \quad i = 2, \ldots, n; k = 2, \ldots, \gamma_i. \end{aligned} \tag{8.11}$$

In addition to the previous notational convention, we now introduce the following further notational convention, which again involves no loss of generality:

(1) The pure-strategy n-tuples a^m of the game have been re-numbered in such a way that the first γ^* pure-strategy n-tuples $a^1, \ldots, a^{\gamma^*}$ will now have the following form. For any m with $1 \leqq m \leqq \gamma^*$, let i and k be the unique pair of numbers satisfying $m^*(i, k) = m$.

Then

$$a^m = \left(a_1^1, \ldots, a_{i-1}^1, a_i^k, a_{i+1}^1, \ldots, a_m^1\right) \tag{8.12}$$

Thus, we can write

$$u_i^m = u^{m^*(i,k)} = U_i\left(a_1^1, \ldots, a_{i-1}^1, a_i^k, a_{i+1}^1, \ldots, a_n^1\right), \text{ for } i = 1, \ldots, n; k = 2, \ldots, \gamma_i. \tag{8.13}$$

Let u^* be the vector formed of those γ^* components u_i^m of vector u which can be written in form (8.13). Let u^{**} be the vector formed of the remaining $(nK - \gamma^*)$ components of u. Hence

$$u = (u^*, u^{**}) \tag{8.14}$$

The set of all possible vectors u^* is a γ^*-dimensional Euclidean space, to be denoted as $\mathscr{I}^* = \{u^*\}$; whereas the set of all possible vectors u^{**} is an $(nK - \gamma^*)$ dimensional Euclidean space, to be denoted as $\mathscr{I}^{**} = \{u^{**}\}$. Clearly, $\mathscr{I}^* \times \mathscr{I}^{**} = \mathscr{I}$.

Since p is an equilibrium point in game Γ, it must satisfy condition (5.2). This condition can also be written as

$$p_i^1 p_i^k \left[U_i\left(a_i^k, \bar{p}_i\right) - U_i\left(a_i^1, \bar{p}_i\right)\right] = 0, \quad \text{for} \quad i = 1, \ldots, n; k = 2, \ldots, \gamma_i. \tag{8.15}$$

Since p_i^1 and $p_i^k > 0, (49)$ is equivalent to (5.2). In view of (3.12), this condition can also be written as

$$\begin{aligned} u^{m^*} = & \sum_{\substack{m \neq m^* \\ m \in M}} \frac{\left[q_i^m\left(a_i^k\right) \prod_{j \neq i} q_j^m\left(\bar{p}_j\right)\right] u_i^m}{\prod_{j \neq i} q_j^{m^*}\left(p_j\right)} \\ & - \sum_{m \in M} \frac{\left[q_i^m\left(a_i^1\right) \prod_{j \neq i} q_j^m\left(\bar{p}_j\right)\right] u_i^m}{\prod_{j \neq i} q_j^{m^*}\left(p_j\right)}, \text{ for } i = 1, \ldots, n; k = 2, \ldots, \gamma_i \end{aligned} \tag{8.16}$$

Here $M = \{1, 2, \ldots, K\}$ and $m^* = m^*(i,k)$. It is permissible to write eq. (5.2) [or (8.15)] in form (8.16) because, by (8.12), we have $q_j^{m^*}\left(p_j\right) = p_j^1$ for all $j \neq i$, and $p_j^1 > 0$ since $a_j^1 \in C\left(p_j\right)$.

Note that each quantity $u_i^{m^*}$ for a specific value of $m^* = m^*(i,k)$ occurs, with a nonzero coefficient, only in one equation of form (8.16) (where it occurs on the left-hand side). This is so because, by (3.10) and (8.12), for any $k' \neq k$, we have $q_i^{m^*}\left(a_i^{k'}\right) = 0$. Therefore, if we know the γ^* components p_i^k of vector π, and know the $(nK - \gamma^*)$ components u_i^m of vector u^{**}, then we can compute each one of the γ^* components $u_i^{m^*}$ of vector u^* separately, from the relevant equation of form (8.16). Consequently, the γ^* equations of form (8.16) define a mapping $\rho : (\pi, u^{**}) \rightarrow u^*$ from set $\Pi \times \mathscr{I}^{**}$ to set $\mathscr{I}^*$. This mapping ρ is continously differentiable because, by (8.6), for each point π in Π we have $p_i^k > 0$ for $k = 2, \ldots, \gamma_i$; so that the denominators on the right-hand side of (8.16) never vanish within Π.

We can use this mapping ρ to define another mapping $\rho^* : (\pi, u^{**}) \rightarrow (u^*, u^{**}) = u$, where $u^* = \rho(\pi, u^{**})$. This mapping ρ^* is from set $\Pi \times \mathscr{I}^*$ to set $\mathscr{I}^* \times \mathscr{I}^{***} = \mathscr{I}$; and it is continuously differentiable since ρ is.

9. Almost All Finite Games are regular: The final Stretch

Theorem 9.1. *Almost all finite games are regular.*

Proof. Instead of using the γ^* equations of form (8.16), we can also use the equivalent γ^* equations of form (8.15), in order to define the mappings ρ and ρ^*. But if we do so, then the Jacobian of mapping ρ^* can be written as

$$J^*\left(\pi, u^{**}\right)=\frac{\partial\left(F_i^2, \ldots, F_i^k, \ldots, F_n^{\gamma_n}\right)}{\partial\left(p_i^2, \ldots, p_i^k, \ldots, p_n^{\gamma_n}\right)}, \quad i=1, \ldots, n; \quad \text{and, for each} \quad k=2, \ldots, \gamma_n,$$

where the F_i^k 's are the functions $F_i^k=F_i^k(0,p)$ defined by (7.4). This means that $J^*\left(\pi, u^{**}\right)$ is a subdeterminant of the Jacobian determinant $J(0,p)$ defined by (7.1, (7.2), and (7.4); it is that particular subdeterminant that we obtain if, for each player i, we cross out the rows and the columns corresponding to $k=1$, and to $k=\gamma_i+1, \ldots, K_i$. It is easy to verify that, owing to the special form of the functions $F_i^1(i=1,\ldots,n)$ as defined by (7.2), and owing to the fact that $p_i^{\gamma_i+1}=\cdots=p_i^{K_i}=0$, this crossing out of these rows and columns does not change the value of the original determinant $J(0,p)$. Hence, $J^*\left(\pi, u^{**}\right)=J(0,p)$ if π is the subvector of p defined by (8.5) and (8.8).

Let $\mathscr{E}\left(C^*\right)$ be the set of all games Γ in $\mathscr{I}$ having at least one irregular equilibrium point p with set $C^*=C(p)$ as its carrier set. Equivalently, $\mathscr{E}\left(C^*\right)$ can also be defined as the set of all vectors $u=\rho^*\left(\pi, u^{**}\right)$ corresponding to those points $\left(\pi, u^{**}\right)$ in set $\left(\Pi \times \mathscr{I}^{**}\right)$ at which the Jacobian $J^*\left(\pi, u^{**}\right)=J(0,p)$ vanishes. By SARD's Theorem [HZ17], this set $\mathscr{E}\left(C^*\right)$ is a set of measure zero in the (nK) dimensional Euclidean space $\mathscr{I}$.

Let $\mathscr{E}^*$ be the set of all games Γ in $\mathscr{I}$ having at least one irregular equilibrium point p, regardless of what its carrier $C^*=C(p)$ is. Thus, $\mathscr{E}^*$ is simply the set of all irregular games in $\mathscr{I}.\mathscr{E}^*$ is the union of a finite number of sets $\mathscr{E}\left(C^*\right)$, corresponding to different carrier sets C^*. Since each set $\mathscr{E}\left(C^*\right)$ is a set of measure zero in $\mathscr{I}$, their union $\mathscr{E}^*$ will also have this property.

Next, we shall show that $\mathscr{E}^*$ is a closed set. Let $\Gamma^1, \Gamma^2, \ldots$ be a sequence of irregular games, with the defining vectors $u^1=u\left(\Gamma^1\right), u^2=u\left(\Gamma^2\right), \ldots$. Suppose that the sequence $u^1, u^2, \ldots$ converges to a given vector u^0. Let Γ^0 be the game corresponding to $u^0=u\left(\Gamma^0\right)$. We have to show that Γ^0 is likewise an irregular game.

Let $p^1, p^2, \ldots$ be a sequence of strategy n-tuples, such that $p^j(j=1,2,\ldots)$ is an irregular equilibrium point in game Γ^j. All these points p^j lie in the compact set P. Consequently, the sequence $\{p^j\}$ must contain a convergent subsequence. Suppose the latter consists of the points $p^{j_1}, p^{j_2}, \ldots$, and that it converges to some point p^0 in P. Then:

(1) This point p^0 will be an equilibrium point of game Λ^0 This is so because the set $Q(\Lambda)$ of all equilibrium points in any given game is an upper semi-continuous set function of the defining vectors $u=u(\Lambda)$ of Λ, i.e., of the payoffs u_i^m of Λ
(2) This point p^0 will be an irregular equilibrium point of game *Lambda* This is so because $J(0,p^{j1})=J(0,p^{j2})=\cdots=0$ since $p^{j1}, p^{i2}, \ldots$ are irregular equilibrium points. Consequently, $J(0,p^0=0$ since p^0 is the limit of the sequence $p^{j1}, p^{j2}, \ldots$ and since $J(0,p)$ is a continuous function of p.

Consequently, p^0 is an irregular equilibrium point in Λ^0 and, therefore, Λ^0 itself is an irregular game, as desired. This completes the proof of Theorem 3.

Theorem 9.2. *In almost all finite games, the number of equilibrium points is finite and odd.*

Proof. From Theorems (7.9), (8.1) and (9.1), this theorem is directly implied.

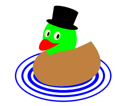

That concludes the proof for Wilson's Oddness Theorem.

10. Some exceptional games

If you've made it this far in this paper, you are clearly interested in the Wilson's Oddness Theorem and Game Theory in general.

This section will have some of the interesting exceptions for Wilson's Oddness Theorem. The theorem itself is a really interesting result. It helps us understand the nature of the number of equilibria being odd and finite in *almost all* situations.

Let's define *almost all* in simpler words.

Imagine you have a dartboard represented in a graph as the unit disk.

If you throw a dart that lands on the dartboard, what's the probability it lands in a particular region?

If you pick the region of the entire disk, then the dart is surely to land in the region so the probability is 1. If you pick a region that is a disk of radius 1/4, then that region comprises 1/8 of the total area, and hence the probability will be 1/8.

To return to Wilson's Oddness Theorem (Theorem 9.2 in this paper), the theorem states that finite games that have an even number of solutions or an infinite number is a set that has measure zero. If you think about the set of finite games as the dartboard, then the games that have an even or infinite number of solutions are like the collection of single points (technical point: a line in a dartboard also has measure zero). Games with an even number of Nash equilibria certainly exist, and the set can even be a collection of an infinite number of items. However, these games are a set of measure zero relative to the entire set of games.

An example of this exception is:

Consider the following game, player x and player y get 1 a piece if they play (Up, Left) and they get nothing otherwise.

Player 'y'

Player 'x'		Left	Right
	Up	1, 1	0, 0
	Down	0, 0	0, 0

Figure 1. Possible Outcomes (Source: Mind Your Decisions)

As shown in blue, Left-Up is one equilibrium point and Right-Down is one equilibrium point. That gives 2 (an even number) of equilibrium points!

In fact, in the Up-Left box, if we change 1 to any positive real number, there will always exist 2 equilibrium points.

There possibly exists an infinite number of these exceptions but they are so few compared to the other games that they are called a close set of measure 0.

11. Further reading

An interesting article [BF21] that can be checked out which applies Wilson's Oddness Theorem is 'Oddness of the number of Nash Equilibria: the case of polynomial payoff function' by Philippe Bich and Julien Fixary.

References

[BF21] Philippe Bich and Julien Fixary. Oddness of the number of Nash equilibria: the case of polynomial payoff functions, 2021. Documents de travail du Centre d'Economie de la Sorbonne 2021.27 - ISSN : 1955-611X.

[Deb70] Gerard Debreu. Economies with a finite set of equilibria. *Econometrica*, 38(3):387–392, 1970.

[Har73] J. C. Harsanyi. Oddness of the number of equilibrium points: A new proof. *International Journal of Game Theory*, 2(1), Dec 1973.

[HZ17] Piotr Hajłasz and Scott Zimmerman. The dubovitskiĭ-sard theorem in sobolev spaces. *Indiana University Mathematics Journal*, 66(2):705–723, 2017.

[Kun97] Ernst Kunz. *Einführung in die algebraische Geometrie*. Jan 1997.

[Nas51] John Nash. Non-cooperative games. *The Annals of Mathematics*, 54(2):286, Sep 1951.

[Wil71] Robert Wilson. Computing equilibria of n-person games. *SIAM Journal on Applied Mathematics*, 21(1):80–87, 1971.