PRINCETON UNIVERSITY

DISCLAIMER NOTICE

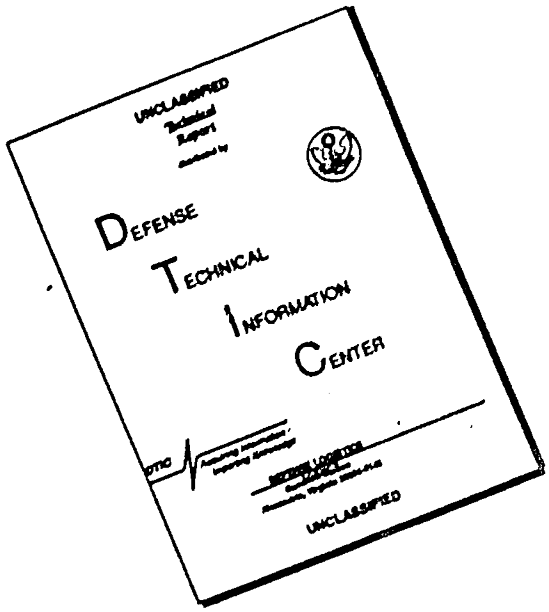

THIS DOCUMENT IS BEST QUALITY AVAILABLE. THE COPY FURNISHED TO DTIC CONTAINED A SIGNIFICANT NUMBER OF PAGES WHICH DO NOT REPRODUCE LEGIBLY.

ON NASH EQUILIBRIUM POINTS AND GAMES OF IMPERFECT INFORMATION

James H. Case
George Kimeldorf

Econometric Research Program
Research Memorandum No. 112

June 1970

The research described in this paper was supported by the Office of Naval Research N00014-67 A-0151-0007, Task No. 047-086.

Princeton University
Econometric Research Program
207 Dickinson Hall
Princeton, New Jersey

ON NASH EQUILIBRIUM POINTS AND GAMES OF IMPERFECT INFORMATION

James Case
George Kimeldorf*

Introduction:

In this paper, we wish to study a very simple class of two player, nonzero-sum games. These will be distinguished by the fact that each player has only local (not global) information regarding the payoff functions, and made simple by the assumption that each player's strategy set be a subset of the real line R. Thus our analysis may be confined to the Euclidean plane, and certain tools will be available to us which have no analogues in higher dimensions. Moreover, we consider no other solution concepts than Nash equilibria.

We restrict ourselves to such simple games, and to a single solution concept, for several reasons. First, the study of non-zero sum games with imperfect information is still in its infancy, and it is our belief that the notion of a Nash equilibrium point is a highly relevant one for such games. For we suspect that, when faced with imperfect information, people and groups of people do in fact attempt to gain additional information, and then to utilize that information in roughly the manner we shall shortly describe. And second, we hope that by confining ourselves to simple games, we shall reveal more of the nature of Nash equilibrium points, the relationships which may exist between them, and the methods by which they may be computed, than is possible by the conventional fixed point procedures.

* Department of Statistics, The Forida State University, Tallahassee, Florida.

1. An Example

We begin with an example. Consider two firms X and Y, which manufacture toothpaste. Suppose that it costs r cents to produce a tube of brand X and s cents for one of brand Y. Let p be the price at which X offers its toothpaste to the public on a given day, and let q by the price for Y. Then if the function Φ be defined for every real number x by the relation

$$(1.1) \qquad \Phi(x) = \int_{-\infty}^{x} e^{-\xi^2}\,d\xi ,$$

so that $0 \leq \Phi(x) \leq \sqrt{\pi}$ and $\Phi(-x) = \sqrt{\pi} - \Phi(x)$, we shall assume that the public demands $m\,\Phi(q-p)$ tubes of brand X during the day and $m\sqrt{\pi} - m\,\Phi(q-p) = m\,\Phi(p-q)$ tubes of brand Y. Here $m\sqrt{\pi}$ is some large positive integer, and is equal to the total number of tubes of toothpaste bought by the public during the day. For simplicity, we have taken this number to be independent of p and q. That is, we have assumed the public's demand for toothpaste to be totally inelastic.[1] But we do <u>not</u> assume that the firms know this!

At the end of the day, the managers of X will know p, because they control it, and their own profit

$$(1.2) \qquad f(p,q) = m(p-r)\,\Phi(q-p) ,$$

because they keep books. Also they will know q, because they will have gone out and bought a tube of brand Y. But they will not

---

[1] It has recently been brought to our attention that there exists a considerable body of statistical data concerning such markets, and that the demand is more likely to take the form: demand for $X = \Phi(\log q/p)$.

know Y's profit

$$(1.3) \qquad g(p,q) = m(q-s)\,\phi(p-q)\ ,$$

or (equivalently) the public's demand for brand Y , because they do not have access to Y's books. In short, X's information at the end of a day is entirely contained in the triple of real numbers $(p,\ q,\ f(p,q))$, while Y's is summarized by the list $(p,\ q,\ g(p,q))$.

But in reality, the firms would never content themselves with so meager a knowledge of the market structure. For by appropriate experiments, they can at least estimate the effects of certain price changes. And for simplicity, we shall allow them to measure exactly, certain quantitites which in practice they could only estimate. To this end, we assume that both brands are sold at an infinite number of outlets (drugstores, supermarkets, etc.) $\theta_1, \theta_2, \ldots$. Then firm X could, on a day when the prices were p and q respectively, select a subsequence $\theta_{n_1}, \theta_{n_2}, \ldots$ of the set of outlets, and reduce the price of brand X at those stores to the levels $p_{n_1}, p_{n_2}, \ldots$, where $\lim_k p_{n_k} = p$. Then X would know the derivative

$$(1.4) \qquad f_p(p,q) = \lim_k \frac{f(p_{n_k},q) - f(p,q)}{p_{n_k} - p}$$

exactly. And Y could measure $g_q(p,q)$ by selecting some different sequence of outlets at which to hold sales. Of course, they could never really hold more than a finite number of sales, at different prices, and in suitably distant locations. But since that number is large, it is felt that $f_p(p,q)$ and $g_q(p,q)$ may be regarded as known. We shall discuss various other experiments, which the firms might wish to perform, a little later on.

In any event, though they do not know it, the firms must, every day, play the game

$$(G_1) \qquad \underset{p \in R}{\text{maximize}} \; f(p,q) \qquad\qquad \underset{q \in R}{\text{maximize}} \; g(p,q) .$$

It was shown in [1] that G_1 has a unique Nash equilibrium point at the intersection of the curves $C = \{(p,q): f_p(p,q) = 0\}$ and $\Gamma = \{(p,q): g_q(p,q) = 0\}$ [1]. C was described geometrically as the graph of a function $p = \varphi(q)$ such that $\varphi'(q) > 0$, $\varphi''(q) > 0$, and $\lim_{q \to -\infty} \varphi(q) = r$. Similarly, Γ is the graph of $\psi(p)$, where $\psi'(p) > 0$, $\psi''(p) > 0$, and $\lim_{p \to -\infty} \psi(p) = s$. Thus C and Γ are as shown in figure 1.

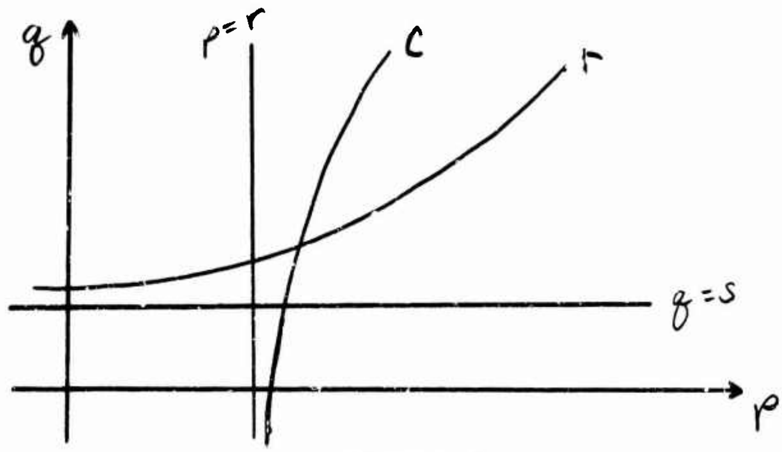

Figure 1.

[1] If we replace $\Phi(p-q)$ by $\Phi(\log p/q)$ in the statement of our problem, the curves C and Γ still have the shape indicated in figure 1. And since the rest of our argument depends only on the shape of c and r, it would be unchanged by the incorporation of the more realistic demand function.

We may describe their relationship by saying that the graph Γ of ψ crosses C from top to bottom as p tends from $-\infty$ to $+\infty$. They do not cross again since neither can meet any line of unit slope more than once. Let P^* be their unique point of intersection, and let (p^*,q^*) be the coordinates of P^*.

We wish to point out that, using only the information that we have allowed them, there are a variety of ways in which the players X and Y may be expected to locate P^*. For instance, on a day when their prices are p_o and q_o respectively, X might select a subsequence $\{\theta_{n_k}\}_{k=1}^{\infty}$ of the set of outlets (in widely separated communities, naturally), and choose a sequence $\{p_{n_k}\}_{k=1}^{\infty}$ of "sale prices" which is dense in the entire real line[1]. In this way, he could locate the exact point (p_1,q_o) at which C crosses the line $q = q_o$, and set his price thereafter at the level $p = p_1$. But having done so, X must surely expect Y to reply in kind, and reset his price at the level $q = q_1$ which is optimal for him against p_1. And if we denote by $P_o, P_1, P_2, \ldots$ the sequence of points whose coordinates are $(p_o,q_o), (p_1,q_o), (p_1,q_1),\ldots,$ it is clear from figure 2 that $\lim_n P_n = P^*$.

---

[1]Or at least in the interval (r,∞).

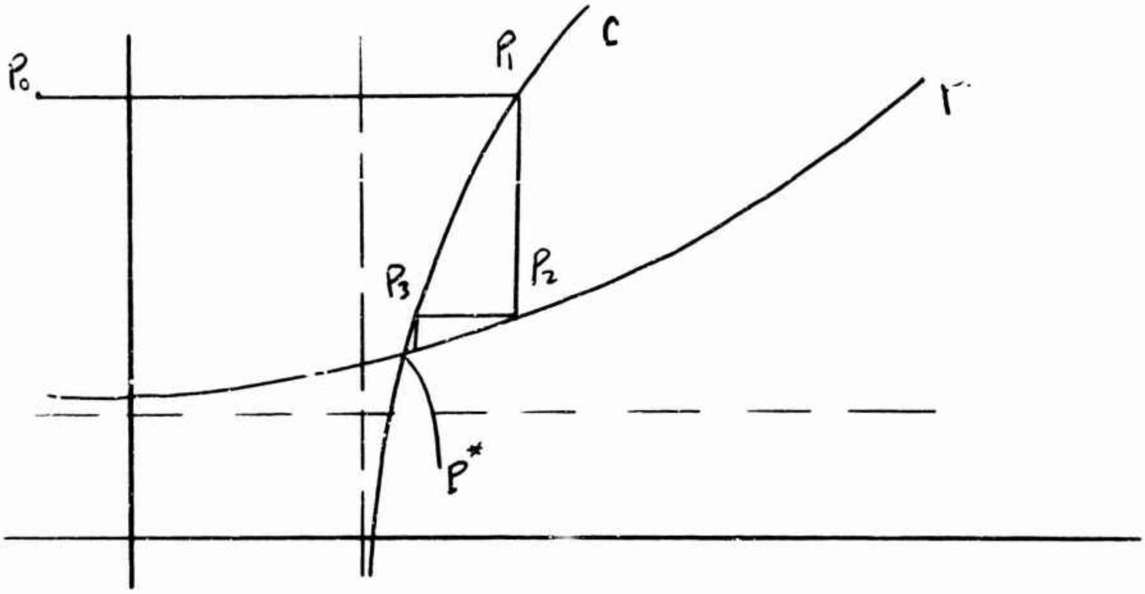

Figure 2

We note too that the result is the same if Y is allowed to make the first move.

Alternatively, X might wish to start, on a day when the firms' respective prices are p_o and q_o, by switching to a new price p_1 which differs from p_o by no more than ϵ. Accordingly, we would choose a sequence $\{p_n\}$ of sale prices dense in the interval $p_o-\epsilon \leq p \leq p_o+\epsilon$, and take p_1 to be the price which optimizes $f(p,q_o)$ thereon. And, not unreasonably, Y might then switch to a price q_1 which maximizes $g(p_1,q)$ over an interval $q_o-\delta \leq q \leq q_o+\delta$. The sequence $P_o, P_1, \ldots$ obtained in this way is indicated in figure 3.

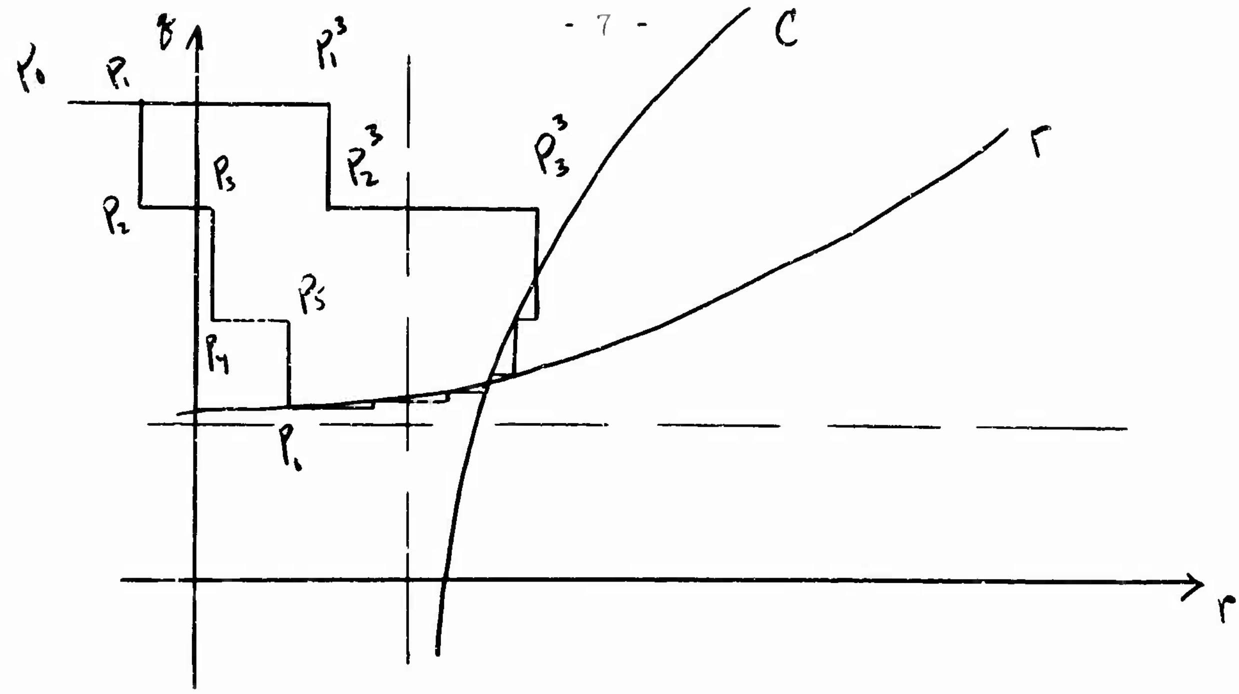

Figure 3

The sequence P_o, P_1^3, P_2^3, ... is the variant of P_o, P_1, P_2, ... obtained by using 3ϵ instead of ϵ . Rather clearly, the process does converge for any pair of positive numbers ϵ and δ . And again, the result is independent of which player is allowed to make the first move. In what follows, we shall refer to the process of figure 2 as the long-step process (π_L) , and that of figure 3 as the short-step process (π_S) . We shall also consider an infintesimal process (π_I) first suggested by Rosen [3]; namely that the players change their prices simultaneously and continuously in time according to the laws

$$
\begin{aligned}
\dot{p}(t) &= f_p(p(t), q(t)) \\
\dot{q}(t) &= q_q(p(t), q(t))
\end{aligned}
\qquad (1.5)
$$

For completeness, we sketch the solutions of the ordinary differential equations (1.5) near P* in figure 4.

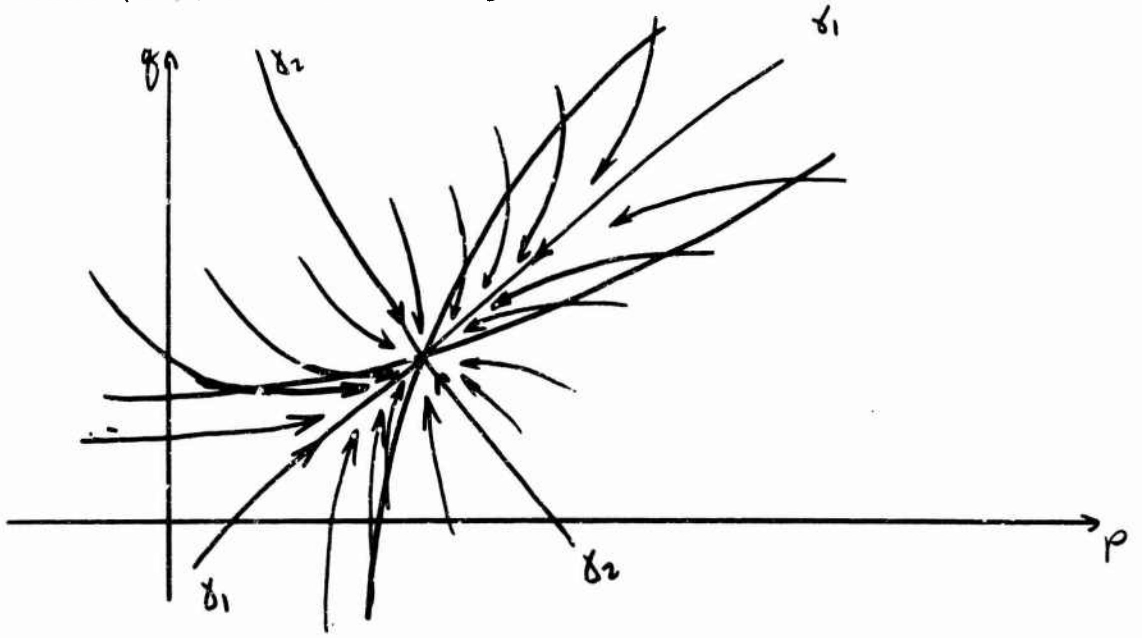

Figure 4

All the solution curves $(p(t),q(t))$, $-\infty < t < \infty$, approach P* as $t \to +\infty$, and all save the one labeled γ_2 are tangent to the single solution curve γ_1. This state of affairs may be summarized by saying that the singularity P* of the system (1.5) is a "stable node". We will have more to say about the singularities of systems like (1.5) later on.

Finally, we observe that if $r=s$, the equations

$$(1.6)\qquad \begin{aligned} -f_p(p,q) &= (p-r)\,\Phi'(q-p) - \Phi(q-p) = 0 \\ -g_q(p,q) &= (q-s)\,\Phi'(q-p) - \Phi(p-q) = 0 \end{aligned},$$

have a symmetric solution $p^* = q^*$, to be found by solving

$$(1.7) \qquad (p-r)\,\Phi'(0) = \sigma(p-r) = \sqrt{\pi}/2 = \Phi(0)$$

for $p^* = r + \sqrt{\pi}/2\sigma = q^*$. And the solution of (1.6) is unique for all values of r and s. Thus at equilibrium, the profit on a tube of toothpaste is $\sqrt{\pi}/2\sigma$ cents. And if we recall that the public's demand for brand X on a day on which $p = q + h$ is $m\,\Phi(h) \doteq m\Phi(0) + m\,h\,\Phi'(0) = m\sqrt{\pi}/2 - m\,h\sigma$, it is clear that $m\sigma$ is just the number of sales lost by firm X if they charge a penny ($h=1$) more for their product than does firm Y. So if the market is highly sensitive to small changes in price, σ is large and profits are small. But if it is insensitive, then profits may be very high indeed.

Of course, if σ is too large, the firms will not likely accept a profit of only $\sqrt{\pi}/2\sigma$ cents a tube. They do not have to because both f and g are increasing functions along both C and Γ. Thus one firm, say X, may choose a price p greater than p^*, the equilibrium price, and allow Y to maximize its own profit against p. Then the firms would find themselves operating at a point P to the right of P^* on Γ, at which they both earn higher profits. And perhaps they can even negotiate a compromise P' above P (but below C) at which their relative shares of the market are more nearly equal. But the answers to such questions go beyond the theory of Nash Equilibria, and we shall not discuss them here. For it is our belief that, confronted with the market described above, firms really do behave much as we have said they would, and arrive at Nash Equilibrium prices.

. Concave Games:

Initially, we shall confine our attention to the class of "concave games" defined on the entire Euclidean plane $R \times R$. Such a game is completely determined once we have specified the payoff functions $f(x,y)$ for player X and $g(x,y)$ for player Y . To play, player X chooses a "strategy" $x \in R$, player Y chooses $y \in R$, and the "bank" (the toothpaste-buying public, in our example) pays $f(x,y)$ dollars to X and $g(x,y)$ dollars to Y .

We shall assume that f and g are at least twice continuously differentiable throughout $R \times R$. Also, we shall assume that $f(x,y)$ is bounded above on each line y = constant, that $f_x(x,y) = 0$ holds at a unique point $(\varphi(y),y)$ on that line, and that $f_{xx}(x,y)$ is negative everywhere. Similarly, $g(x,y)$ must be bounded above on point $(x, \varphi(x))$ on that line, and $g_{yy}(x,y)$ must be negative everywhere. It then follows, from $g_y(x, \psi(x)) \equiv 0$, that $\psi'(x) = -g_{xy}(x, \psi(x))/g_{yy}(x, \psi(x))$ is well defined and continuous on the whole real line, and that because the denominator is never zero, the graph of the function ψ can never be tangent to any line x = const. Similarly, the graph of φ can never be tangent to a line y = const.

A strategy x_o for X is said to be "rational for X" against a particular strategy y_o for Y if $f(x_o,y_o) > f(x,x_o)$ for every $x \neq x_o$. And the curve $C = \{(x,y): f_x(x,y) = 0\}$ is called the "rational curve" for X . Similarly, $\Gamma = \{(x,y): g_y(x,y) = 0\}$ is called the rational curve for Y . Any point (x,y) in their intersection is rational for both, and is said to be a Nash equilibrium point of the game

$$\text{(G)} \qquad \underset{x\in R}{\text{maximize}}\ f(x,y) \qquad\qquad \underset{y\in R}{\text{Maximize}}\ g(x,y) \ .$$

That is, if (x_o,y_o) is in both C and Γ , then

$$\text{(2.1)} \qquad f(x_o,y_o) \geq f(x,y_o) \qquad \text{and} \qquad g(x_o,y_o) \geq g(x_o,y)$$

for every other point (x,y) in RxR. We assume that C and Γ meet at only a finite number of points.

They need not meet at all. For if (example 1) $f = xy - \frac{1}{2}x^2$ and $g = (x+1)y - \frac{1}{2}y^2$, then $f_{xx} = -1 = g_{yy}$, C is the line $x = y$, and Γ is the line $y = x+1$. On the other hand, they may meet any finite number of times. For if (example 2), f is again equal to $xy - \frac{1}{2}x^2$, and $g = (x+p(x))y - \frac{1}{2}y^2$, then $f_{xx} = -1 = g_{yy}$ as before, C is still the line $x = y$, and Γ becomes the graph of the polynomial $x+p(x)$. So if $p(x) =$ $= (x-1)(x-2)\ldots(x-n)$, Γ crosses C at each of the points $(1,1), (2,2),\ldots,(n,n)$, and nowhere else. The interesting questions in the theory of planar concave games concern not the location of the Nash equilibria, but the various types of stability which they may have.

3 Finite Stability of the Nash Equilibria

We shall consider the long-step process π_L first. The process $\pi_L(X)$ starting from the point P_0, whose coordinates in RxR are x_0, y_0, and in which player X is allowed to move first, may be defined inductively as follows: $\pi_L(X)$ is that sequence $P_0, P_1, \ldots$ of points in the plane (having the coordinates $(x_0, y_0), (x_1, y_1), \ldots$) for which P_{2k+1} is the (unique) intersection of the line $y = y_{2k}$ with C, for every $k=0, 1, \ldots$, and for which P_{2k+2} is the (also unique) intersection of the line $x = x_{2k+1}$ with Γ. The process $\pi_L(Y)$ is defined analogously.

It is clear from our example in §1, that both $\pi_L(X)$ and $\pi_L(Y)$ may converge, from every starting point (x_0, y_0) in RxR to the same equilibrium point (x^*, y^*) of the game G. But this need not be the case. For if (example 3) the polynomial $p(x)$ in example 2 were just $p(x) = x$, Γ would be the line $y = 2x$ and, as indicated in figure 5, the sequences π_L all diverge to ∞.

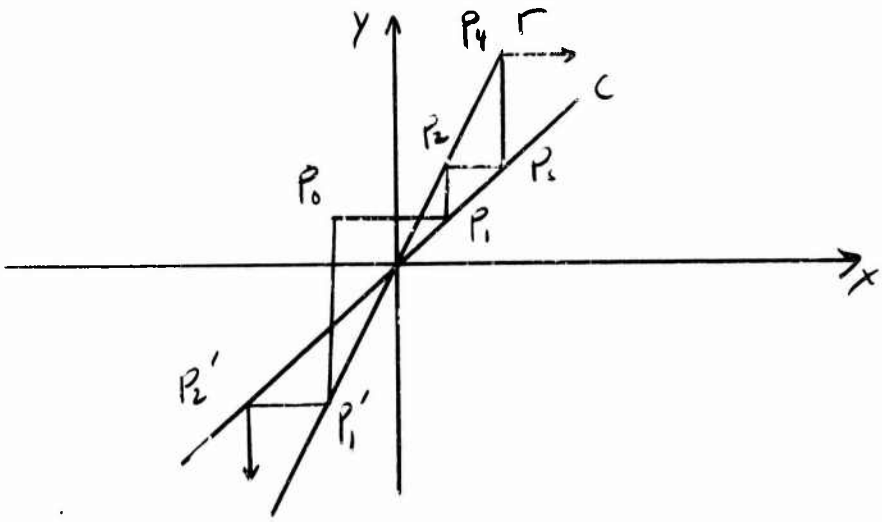

Figure 5

Here the points $P_o, P_1, P_2, \ldots$ form the sequence $\pi_L(X)$, while $P_o, P_1', P_2', \ldots$ form $\pi_L(Y)$. We note that the two sequences diverge not only from the unique Nash equilibrium point (0,0) of the game, but also from each other. Indeed direct calculation reveals that, for the game defined by the payoff functions

$$(3.1) \qquad f = \frac{1}{2} a x^2 + bxy - F(y) \quad \text{and} \quad g = \frac{1}{2} \alpha y^2 + \beta xy - G(x) ,$$

where $a = f_{xx}$ and $\alpha = g_{yy}$ are negative numbers and F and G are arbitrary functions, the sequences $\pi_L(X)$ and $\pi_L(Y)$ both converge to the origin from any $P_o \neq (0,0)$, whenever the ratio $b\beta/a\alpha$ is less than one in absolute value (i.e., $|b\beta| < a\alpha$), and diverge to ∞ if it is greater. If $b\beta/a\alpha = -1$, the sequences are both periodic, since $P_o = P_4 = P_4'$, for every starting point P_o. If $b\beta/a\alpha = +1$, Γ and C coincide.

The situation is more complicated still if we consider functions f and g which lead to curved paths C and Γ. For instance (example 3) if

$$(3.2) \qquad f = xy^{1/3} - (1/2)\, x^2 \qquad g = (x^3 - 3x)y - y^2 ,$$

C and Γ are the graphs of $x = \varphi(y) = y^{1/3}$ and $y = \psi(x) = (x^3 - 3)/2$ respectively. Then, as indicated in Figure 6, the square with corners $x = \pm 1$, $y = \pm 1$ is invariant under the process π_L, and the process paths starting both within and without the square converge to it.

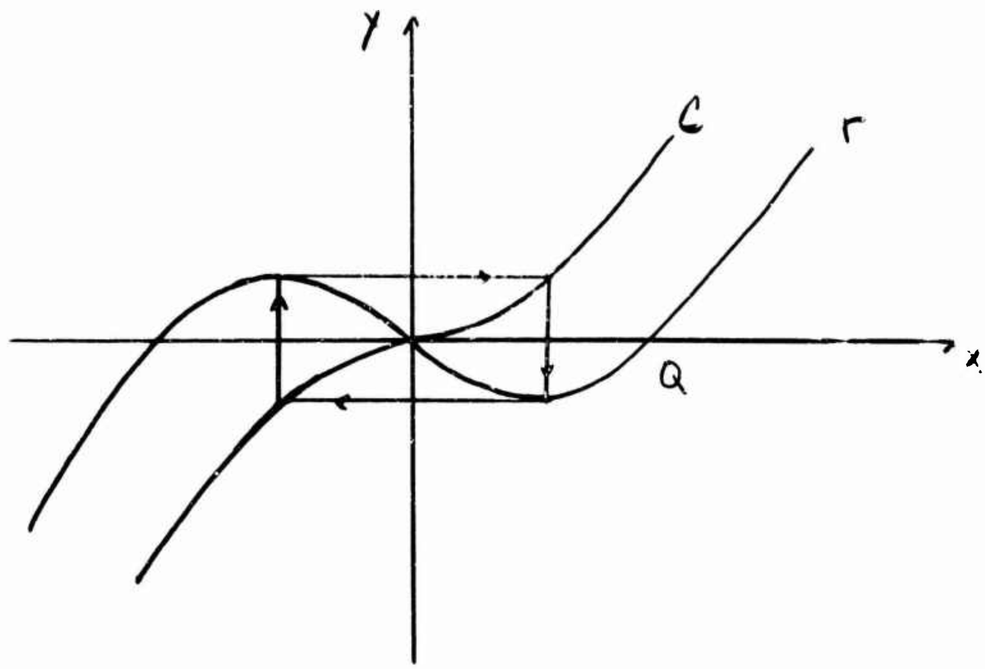

Figure 6

Finally, in games with multiple equilibria, the process π_L may be convergent but ambiguous, in the sense that, from certain starting points P_o, $\pi_L(X)$ and $\pi_L(Y)$ may converge to different equilibrium points P^* and P^{**}. Thus if (example 4)

$$(3.3) \qquad f = xy - (1/2)x^2 \qquad \text{and} \qquad g = x^{1/3}\,y - (1/2)y^2 \;,$$

the points $P^* = (1,1)$ and $P^{**} = (1,-1)$, as well as the origin, are Nash equilibria. And, as an examination of Figure 7 will reveal, $\pi_L(X)$ always converges to P^* from starting points P_o

in the second and fourth quadrants, while $\pi_L(Y)$ goes to P^{**}. Both procedures lead to P^* if P_O is in the first quadrant and to P^{**} if P_O is in the third. The origin is an unstable equilibrium in the sense that, only if P_O lies on a coordinate axis, does even one of the sequences $\pi_L(X)$ and $\pi_L(Y)$ lead there.

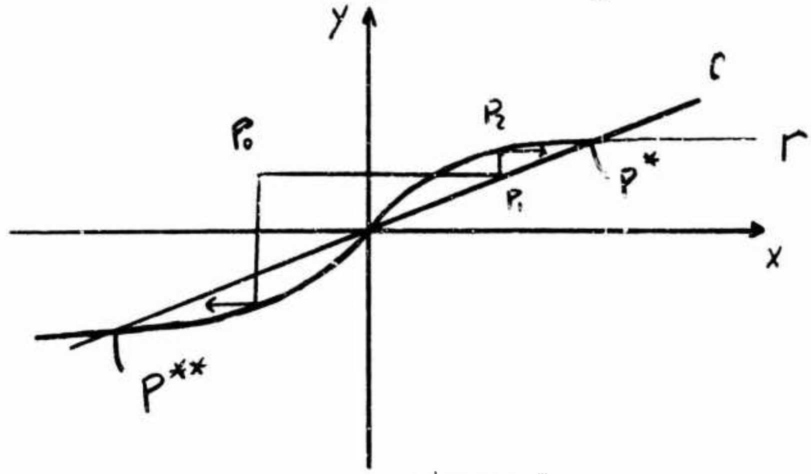

Figure 7.

To fix ideas we say an equilibrium point P^* is L-stable if there exists a neighborhood N of P^* such that starting from any $P_O \epsilon N$ both $\pi_L(X)$ and $\pi_L(Y)$ converge to P^*. Without loss of generality, we may take the equilibrium point P^* under consideration to be the origin. We shall show below that O is L-stable if for some $\varepsilon > O$

(3.4) $$|\varphi(\psi(x))| < |x| \qquad \text{for} \qquad O < |x| < \varepsilon$$

or equivalently, if for some $\varepsilon' > O$

(3.5) $$|\psi(\varphi(y))| < |y| \qquad \text{for} \qquad O < |y| < \varepsilon'.$$

In particular, O is L-stable if $|\varphi'(O)\psi'(O)| < 1$ and is L-unstable if $|\varphi'(O)\psi'(O)| < 1$. Geometrically, this condition states that O is L-stable if the slope of Γ at O is less in absolute value than that of C, and is L-unstable if it is greater. If C and Γ are straight lines the neighborhood N can be taken as the entire plane.

With curved rational curves C and Γ, there exist isolated points from which $\pi_L(X)$ or $\pi_L(Y)$ converge to L-unstable equilibria in a finite number of steps. For example, in Figure 6 starting from the point Q at which Γ intersects the positive x axis, $\pi_L(X)$ converges to the origin in one step and $\pi_L(Y)$ in two steps. Hence starting from any point Q_0 from which Q is reached in a finite number of steps, π_L converges to the origin. Moreover, the set of all Q_0 is unbounded.

In studying the concept of L-stability in a global sense, we shall want to exclude such isolated points. Therefore, given that O is an L-stable equilibrium we seek the largest set M of points from which both $\pi_L(X)$ and $\pi_L(Y)$ converge to O through a path which lies entirely in M; that is, each line segment joining P_i and P_{i+1} must lie in M. We shall call M the <u>convergence region</u> of π_L. The following theorem characterizes the convergence region M.

<u>THEOREM:</u> Let O be an L-stable equilibrium. Then there exists a unique open rectagle M for which:

(a) $\pi_L(X)$ and $\pi_L(Y)$ converge to O from any point $P_0 \in M$ through a path lying wholly in M, and

(b) From any finite vertex of M either

(1) $\pi_L(X)$ and $\pi_L(Y)$ both cycle, or

(2) $\pi_L(X)$ and $\pi_L(Y)$ both converge in at most three steps to an L-unstable equilibrium.

Moreover, M is the largest set having property (a). In particular, M is the rectangle determined by (x_0, y_0), (x_0, y_1) (x_1, y_1) and (x_1, y_1) and (x_1, y_0) where

(3.6)
$$
\begin{aligned}
x_0 &= \sup\{x<0:\ \varphi(\psi(\varphi(\psi(x)))) = x \ \text{ or } \ \varphi(\psi(\varphi(\psi(x)))) = \varphi(\psi(x)) \neq 0\} \\
x_1 &= \inf\{x>0:\ \varphi(\psi(\varphi(\psi(x)))) = x \ \text{ or } \ \varphi(\psi(\varphi(\psi(x)))) = \varphi(\psi(x)) \neq 0\} \\
y_0 &= \sup\{y<0:\ \psi(\varphi(\psi(\varphi(y)))) = y \ \text{ or } \ \psi(\varphi(\psi(\varphi(y)))) = \psi(\varphi(y)) \neq 0 \\
y_1 &= \sup\{y>0:\ \psi(\varphi(\psi(\varphi(y)))) = y \ \text{ or } \ \psi(\varphi(\psi(\varphi(y)))) = \psi(\varphi(y)) \neq 0
\end{aligned}
$$

The proof of the thoerem involves the following lemma whose proof appears in the Appendix.

<u>LEMMA:</u> Let h be a continuous function for which $h(0) = 0$ and for some $\varepsilon > 0$

(3.7)
$$0 < |t| < \varepsilon \implies |h(t)| < |t| .$$

Let t_0 be the largest negative t and t_1 be the smallest positive t for which either $h(h(t)) = t$ or $h(h(t)) = h(t) \neq 0$. Let $I^- = (t_0, 0)$, $I^+ = (0, t_1)$, and $I = (t_0, t_1)$. If $h^{[n]}$ represents the n-fold composition of h with itself, then

(i) $t \in I^+ \implies h^{[n]}(t) < t$, $\quad t \in I^- \implies h^{[n]}(t) > t$,

(ii) h maps the pair $\{t_0, t_1\}$ into itself,

(iii) $t \in I \implies h^{[n]} \in I$,

(iv) $t \in I \implies \lim_{n\to\infty} h^{[n]}(t) = 0$.

To prove the theorem we note first that the successive points of $\pi_L(Y)$ from $P_o = (x,y)$ are $P_1 = (x,\psi(x))$, $P_2 = (\varphi(\psi(x)), \psi(x))$ $P_3 = (\varphi(\psi(x)), \psi(\varphi(\psi(x))))$, . . . Hence (3.4) or, equivalently, (3.5) is necessary and sufficient for 0 to be L-stable. Now let M be the open rectangle of the theorem and $P_o = (x,y) \epsilon M$. That the abscissas of P_i belong to (x_o,x_1) and converge to 0 follows by applying parts (iii) and (iv) of the lemma with $h = \varphi(\psi)$. (We note that (3.4) and (3.5) imply that "sup" and "inf" in (3.6) can be replaced by "max" and "min" respectively, and hence always exist, although they may be infinite.) To show the ordinates belong to (y_o,y_1) we first show that $y_o < \varphi(x) < y_1$; in fact for every $x \epsilon (x_o,x_1)$ we must have $\psi(x) \epsilon (y_o,y_1)$. For suppose $\psi(x) = y_i$ for i=0 or 1 so that $\psi(\varphi(\psi(\varphi(\psi(x))))) = \psi(x)$ or $\psi(\varphi(\psi(\varphi(\psi(x))))) = \psi(\varphi(\psi(x))) \neq 0$. Then $\varphi(\psi(\varphi(\psi(\varphi(x))))) = \varphi(\psi(x))$ or $\varphi(\psi(\varphi(\psi(\varphi(\psi(x)))))) = \varphi(\psi(\varphi(\psi(x))))$. Hence either $h(x) \notin I$ or $h(x) = 0$. The latter possibility can be eliminated since $\varphi(y_i) = 0$ contradicts the definition of y_i . Therefore, by part (iii) of the lemma, $x \notin I$. Applying the lemma with $h = \psi(\varphi)$ we have that the ordinates of P_i belong to (y_o,y_1) and converge to 0 . A similar argument holds for $\pi_L(X)$ and part (a) is proved. The proof of part (b) is elementary. That M is the largest such rectable follows from part (b) since a continuous path from any point outside M to the origin must intersect the boundary.

An example in which (b-1) obtains is the game

$$(3.8) \qquad f(x,y) = xy - x^2 \qquad\qquad g(x,y) = -y^2 - xy - x^3y$$

in which C is the graph of $x = \varphi(y) = y$ and Γ is the graph of $y = \psi(x) = \frac{-x}{2}(1+x^2)$. The origin is L-stable and the convergence region of π_L is the open rectangle with vertices $(\pm 1, \pm 1)$. (b-2) obtains in any game for which φ and ψ are either both increasing or both decreasing. This situation is exemplified by the game (3.3) and illustrated in Figure 7. The points P* and P** are L-stable equilibria with corresponding convergence regions M* , the first quadrant, and M** , the third quadrant. In the second and fourth quadrants π_L is ambiguous. If φ and ψ are either both increasing or both decreasing and there are multiple equilibria at which C and Γ actually cross, then moving along either curve the L-stable and L-unstable equilibria alternate and the L-unstable equilibria partition the plane into a "checkerboard" of rectables; each "black" rectangle is the π_L convergence region of an L-stable equilibrium, while from the "white" rectangles π_L is ambiguous.

The short-step processes π_s have rather more satisfactory convergence characteristics, both in the sense that they sometimes converge when the corresponding long-step process does not, and in that they reduce the size of the set of ambiguous starting points P_o.

A particular short-step process $\pi_s(X; \epsilon, \delta)$ may be defined inductively as follows: $\pi_s(X; \epsilon, \delta)$ is that sequence of points $P_o, P_1, \ldots$ (with coordinates $(x_o, y_o), (x_1, y_1), \ldots$) in RxR for which x_{2k+1} is the (unique) solution of the problem

(3.9) $$\underset{x}{\text{maximize}} \quad f(x, y_{2k})$$

$$\text{subject to} \quad x_{2k} - \varepsilon \leq x \leq x_{2k} + \varepsilon ,$$

y_{2k+1} is the (also unique) solution of

(3.10) $$\underset{y}{\text{maximize}} \quad g(x_{2k+1}, y)$$

$$\text{subject to} \quad y_{2k+1} - \delta \leq y \leq y_{2k+1} + \delta ,$$

and for which $x_{2k+2} = x_{2k+1}$ and $y_{2k+1} = y_{2k}$ for each $k=0,1,\ldots$ $\pi_S(Y;\varepsilon,\delta)$ may be similarly defined. Now let L_{2k} be the horizontal line segment of lenth 2ε centered at P_{2k}. If L_{2k} contains a point of C, then that point is P_{2k+1}. Otherwise P_{2k+1} is the endpoint of L_{2k} which is closer to C. And if L_{2k+1} is the vertical line segment of length 2δ about P_{2k+1}, P_{2k+2} is the point on L_{2k+1} nearest to Γ.

Locally π_S and π_L are identical in that P^* is L-stable if and only if it is S-stable. And if, as in the games (3.1), C and Γ are straight lines, then local stability is equivalent to global stability for either process. Indeed, if we define a convergence region for π_S as we did for π_L (a set M of points from which the process converges to a given stable equilibrium thru a path lying entirely in M), it is clear that the convergence regions would be identical.

Hence, for fixed ε and δ, let us define the <u>convergence region</u> M_S of π_S with respect to a given S-stable equilibrium P^*

to be merely the set of points P_0 from which the processes $\pi_S(X;\cdot,\cdot)$ and $\pi_S(Y;\cdot,\cdot)$ both converge to P^*. Then if φ and ψ are both increasing or both decreasing, the convergence region M_S of π_S contains the convergence region M_L of π_L. For in M_L, the paths of π_S are monotone toward 0, and cannot end at any point other than an intersection of Γ and C. But M_S may be much larger than M_L. For instance if, in the game (3.3) illustrated by Figure 7, we choose $\epsilon = \delta$ very small, the region M_S corresponding to P^* is approximately the half-plane $x + y > 0$, while that corresponding to P^{**} nearly fills the region $x + y < 0$.

4. Infinitessimal Stability.

The process π_I starting from the point P_0, whose coordinates in RxR are (x_0, y_0), may be defined to be the solution $x(t), y(t)$; $t \leq 0$, of the ordinary differential equations

$$\text{(4.1)} \quad \begin{aligned} \dot{x}(t) &= f_x(x(t), y(t)) \\ \dot{y}(t) &= g_y(x(t), y(t)) \end{aligned}$$

for which $x(0) = x_0$ and $y(0) = y_0$. Clearly, since we assumed f and g to be at least twice continuously differentiable, there is only one such solution. And if

$$\lim_{t\to\infty} x(t) = x^* \qquad \text{and} \qquad \lim_{t\to\infty} y(t) = y^*$$

for every pair (x_0, y_0) in some neighborhood of (x^*, y^*), we shall say that the equilibrium point (x^*, y^*) of the game G is "I-stable".

For the games (4.1), the equations (4.1) reduce to the linear equations

$$(4.3) \qquad \dot{x} = ax + by$$
$$\dot{y} = \beta x + \alpha y$$

so that the origin O is the only equilibrium point, and it is I-stable if and only if both roots of the characteristic equation (see , pp.)

$$(4.4) \qquad f(\lambda) = \begin{vmatrix} a-\lambda & b \\ \beta & \alpha-\lambda \end{vmatrix} = \lambda^2 - (a+\alpha)\lambda + (a\alpha - b\beta) = 0$$

have negative real parts. But the roots of (4.4) are just

$$(4.5) \qquad \lambda = \frac{1}{2}\left(a+\alpha \pm \sqrt{(a-\alpha)^2 + 4b\beta}\right),$$

where the radical denotes, as usual, the positive square root. So if the quantity under the radical is not positive, both roots have the real part $(a+\alpha)/2$, which is negative because $f_{xx} = a$ and $g_{yy} = \alpha$ are. And if that quantity is positive, both roots are real, the smaller one is always negative, and the larger one is also negative when $a\alpha > b\beta$, zero when $a\alpha - b\beta = 0$, or positive when $a\alpha < b\beta$. In short, O is I-stable if $a\alpha > b\beta$, or equivalently, iff $\varphi'(0)\psi'(0) < 1$. Geometrically, I-stability requires either that C and Γ lie in different quadrants or that the slope of Γ be smaller in absolute value than that of C. If $\Delta = a\alpha - b\beta = 0$, C and Γ coincide and there are infinitely many neutrally-stable equilibria, a possibility we do not wish to discuss.

The origin may be I-stable for the games (3.1) without being either L-stable or S-stable (as is the case, for example, when $a = \alpha =$ and $b = -\beta = 2$). But L-stability and S-stability always imply I-stability, since $b\beta \leq |b\beta|$. This is one reason for our conviction that I-stability is a more fundamental concept than the other two. We point out that if O is I-unstable, it is a "saddle point" of the system (4.3), and the solution curves have the form indicated in figure 8.

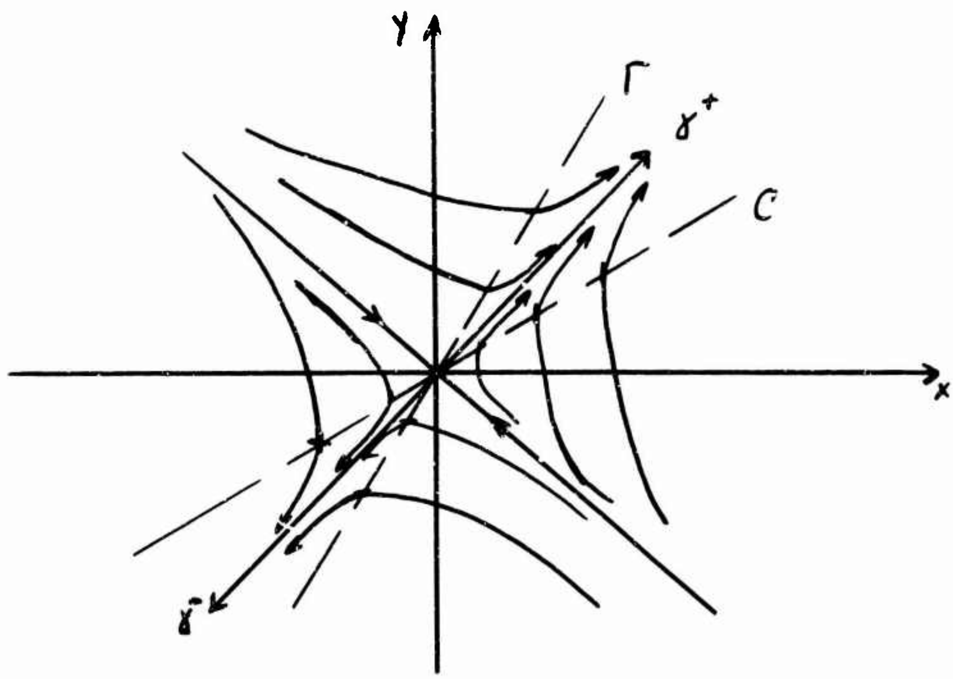

Figure 8.

The rays γ^+ and γ^- are straight-line solutions of (4.3) which are directed away from the origin, and are called (see [2], pg. 212) separatrices. The other pair of straight-line solutions are called separatrices also, and are directed toward the origin. Observe that the lines C and Γ divide the plane into four quadrants, and that each quadrant contains exactly one separatrix. In particular, in the case we have shown, C and Γ have positive slopes, so the outwarddirected separatrices [illegible] lie in the first and third quadrants, respectively. We remark that the distinguished solutions z_1 [illegible] are also called separatrices, and are all directed toward [illegible] the origin. Separatrices play an important role in the analysis of two-dimensional systems of ordinary differential equations and are discussed extensively by Lefschetz in chapters X and IX of [illegible]

If O is [illegible]stable, it may be either a stable node (two negative real roots) as shown in Figure [illegible], or a spiral point (complex roots). In the latter case the solutions of [illegible] are logarithmic spirals which approach the origin asymptotically as t becomes infinite.

A few simple cases deserve mention. If, for instance, we relax our hypotheses temporarily to allow a, [illegible], F, and G all to vanish, [illegible] 1 becomes the simplest example of a non-zero-sum game with bilinear payoffs. The characteristic equation [illegible] reduces to $\lambda^2 = b\beta$, so that the origin may be either a saddle point (if $b\beta > 0$) or a "center" (if $b\beta < 0$). In the latter case, the solutions of [illegible] are ellipses centered at O, which now has only neutral stability. Or if $\beta = -b$, the games [illegible] 1 are essentially zero-sum. For then the equations (4.5) are of the form

$$\dot{x} = f_x(x,y) \quad \text{and} \quad \dot{y} = -f_y(x,y),$$

and their solutions may be expected to lead to a minimax point of the game

$$\max_x \min_y f(x,y).$$

where f is just ax^2 [illegible] bx [illegible] y [illegible]. Indeed, this expectation is realized, because $\beta = -b$ implies that the characteristic function (4.4) is positive when $\lambda = 0$ and roots must both be either real and negative or complex with negative real parts. In either case the origin is a stable equilibrium point for the system (4.3).

Finally, if $a = 0 = \alpha$ and $\beta = 1 = -b$, we are left with the classical penny matching game, whose matrix is

y \ x	-1	-1
-1	-1	-1
1	1	-1

At each play of the game, the players X and Y decide whether to play heads or tails. Then they reveal their choices to the referee. If they both have made the same choice, he awards X's penny to Y, but if they have chosen differently, he gives Y's penny to X. Here x and y are strategy mixtures for X and Y, the values 1 and -1 corresponding to the pure strategies "heads" and "tails" respectively. And the game has the single minimax pair $x = 0$, $y = 0$, at which both players play either heads or tails with equal probability.

But our players would never discover this by the procedure π_I. For the solutions of (4.3) are now given by

$$\begin{pmatrix} x(t) \\ y(t) \end{pmatrix} = \begin{pmatrix} \cos t & -\sin t \\ \sin t & \cos t \end{pmatrix} \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} \tag{4.9}$$

and describe circles about the origin in a counter clockwise direction. This corresponds to the well known fact from ordinary game theory, that the method of repeated (or 'ficticious') play need not converge if the contestants remember only the most recent play of the game. For then the loser always changes to his other pure strategy before the next play, while the winner stands pat, and the successive strategy pairs $(x_0, y_0), (x_1, y_1), \ldots$ cycle about from corner to corner of the unit square again in a counter clockwise fashion. Traditionally, one avoids this difficulty by having recourse to the Brown-Robinson technique, whereby each player recalls all previous plays of the game, and assumes that his opponent will choose, in the upcoming play, that pure strategy which he has chosen most often in the past. But for nonzero-sum (bimatrix) games, even the Brown-Robinson technique does not work. So it seems doubly remarkable that augmenting concavity (blinearity) with the slightly stronger conditions $a < 0$ and $\alpha < 0$ should so alter the problem as to render even the much less sophisticated procedure π_I effective.

Moreover this fact does not depend on the dimension of the strategy spaces. For Rosen [3] has shown that every "strictly diagonally concave" game has a unique Nash equilibrium, and that the process we have called π_I converges to it. The class of diagonally strictly concave games includes all zero-sum games (4.7) for which the n-vector x must be chosen from a compact convex set S in E_n, y from a compact convex Σ in E_m, and for which the Jacobian matrices Φ_{xx} and Φ_{yy} are negative and positive definite respectively. Many other games are included as well. But we do not wish to digress further from the subject of concave games in the plane.

If the functions f and g are not polynomials of order two, the equations (4.1) are not linear, and the curves C and Γ are no longer straight lines. So there may be many Nash equilibria. However, near any one of them, the solutions of (4.1) behave much as though the equations were linear. To see this, we shall make use of Poincare's theory of two dimensional systems. This theory has been considerably extended and refined by Bendixson and others, and is admirably exposed in Lefschetz' book [2].

Let (x_0,y_0) be a point of $C \cap \Gamma$. Then in terms of the new variables $\xi = x-x_0$ and $\eta = y-y_0$, the equations (4.1) may be rewritten

$$(4.10) \qquad \begin{aligned} \dot{\xi} &= a\,\xi + b\eta + f^*(\xi,\eta) \\ \dot{\eta} &= \beta\,\xi + \alpha\eta + g^*(\xi,\eta) \end{aligned} ,$$

providing that f and g have at least continuous third order partial derivatives near (x_0,y_0). Here $a = f_{xx}(x_0,y_0)$, $b = f_{xy}(x_0,y_0)$, $\beta = g_{xy}(x_0,y_0)$, and $\alpha = g_{yy}(x_0,y_0)$, so that $a < 0$ and $\alpha < 0$. The right hand sides of (4.10) are the leading terms of the Taylor expansions of f and g about (x_0,y_0). The functions f^* and g^* are the remainder terms in those expansions, and satisfy

$$(4.11) \qquad \lim_{r \to 0} \frac{f^*(\xi,\eta)}{r} = 0 \quad \text{and} \quad \lim_{r \to 0} \frac{g^*(\xi,\eta)}{r} = 0 ,$$

where $r = (\xi^2 + \eta^2)^{\frac{1}{2}}$ is the distance from (ξ,η) to (x_0,y_0). The constant terms in the expansions do not appear because f_x and g_y vanish at (x_0,y_0).

Under the above assumptions (namely that f and g are of class C^3), we may conclude (see [2], pg. 177), that the remainder terms have

have but a negligible effect in the immediate vicinity of (x_o,y_o), and hence that (x_o,y_o) is a saddle point of the system (4.1) if the characteristic equation (4.4) has a positive real root, a stable node if there are two negative real roots, and a spiral point if both roots are complex. Or, since (a,b) and (β,α) are normal vectors to the curves C and Γ at (x_o,y_o), we may summarize by saying that the equilibrium at (x_o,y_o) is stable if the angle θ between (a,b) and (β,α) is negative, and unstable (a saddle point) if θ is positive.

But if C and Γ are no longer straight lines, it is also possible for them to meet tangentially. In this case θ is zero, and one of the roots of (4.4) vanishes. Then the other root must be $\lambda = a+\alpha < 0$. So, by a theorem of Bendixson, (see [2], pg. 230) we may conclude that a point (x_o,y_o) at which C and Γ meet and share a common tangent is either a node, a saddle point, or a third configuration consisting (in Lefschetz' terminology) of "two hyperbolic sectors and a fan". The latter configuration is sketched below.

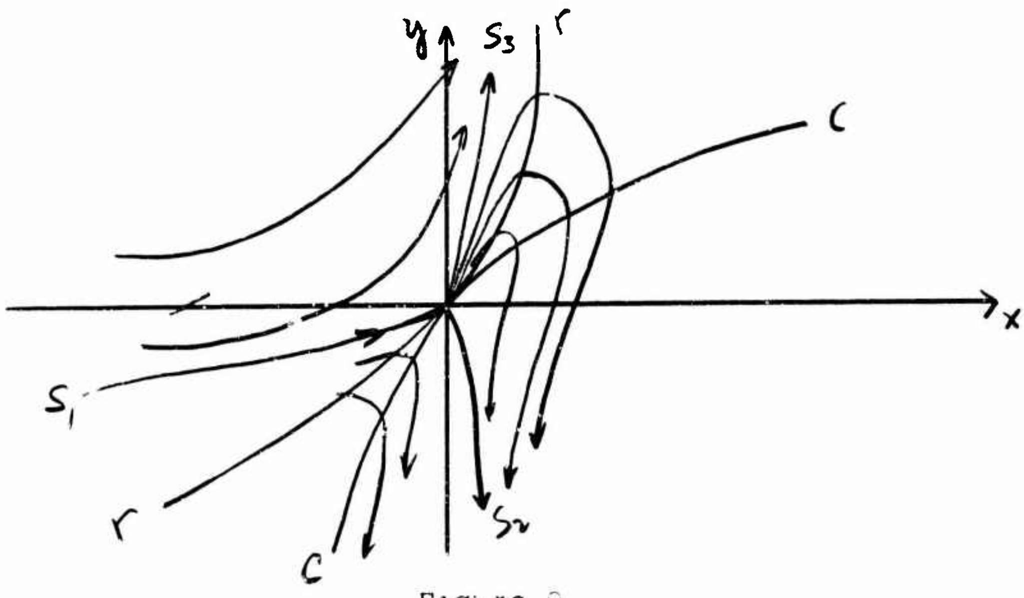

Figure 9

The two "hyperbolic sectors" are the regions immediately above and below the distinguished solution S_1, which are filled with curves which approach either S_2 or S_3 as t becomes infinite. The "fan" is the region across $S_2 \cup S_3$ from S_1, so named because all the solutions in it eminate from the origin (they are all tangent to C and Γ, originally), and "fan out" to fill the right half plane. Moreover, the "indices" of the three types of equilibrium are 1, -1, and 0 so that the index serves to distinguish the various possibilities.

The index of an isolated singularity (equilibrium) (x_0, y_0) of the system (4.1) about a particular oriented rectifiable Jordan curve J surrounding (x_0, y_0) is defined by

$$(4.12) \qquad \text{Index } J = \frac{1}{2\pi} \oint_J \frac{fdg - gdf}{f^2 + g^2} = \frac{1}{2\pi} \oint_J \text{arc tan} \frac{g}{f} ,$$

and has the value $(p-n)/2$, where p is the number of times the field vector $f_x(x(t), y(t)), g_y(x(t), y(t))$ crosses the (upward) vertical ray thru $(x(t), y(t))$, as the point $(x(t), y(t))$ describes the curve J once in the positive direction, and n is the number of times it crosses in the negative direction. It is clear from the expression (4.12) that Index J is a continuous function of the curve J, which takes on only integer values, and hence must be independent of J.

It is easy too, to show (see [2], pg. 186) that if the determinant $\Delta = a\alpha - b\beta$ is not zero, the index of the equilibrium point (x_0, y_0) of the system (4.10) is the same as it would be if f^* and g^* were identically zero, that is 1 for a node or a spiral point and -1 for a saddlepoint. Thus, finally, we may determine the nature stability of a given equilibrium point from the nearby behavior of the curves C and Γ alone. For if Γ crosses C from top to bottom in the direction of increasing x at (x_0, y_0), the index of (x_0, y_0) is 1 and the equilibrium is stable (either a node or a spiral point). Or if Γ crosses C from bottom to top, the index is -1 and the equilibrium is unstable (a saddle point). And if Γ does not cross C but only strikes it tangentially, the index is zero and the equilibrium is again unstable (two hyperbolic sectors and a fan).

This completes our discussion of local phenomena. We have shown that the behavior of the system (4.1) near an isolated equilibrium is completely determined by the manner in which the curves C and Γ intersect there. In the next section, we shall consider the multiplicity of equilibria and discuss the global behavior of the system (4.1).

5. The Global Configuration

In order to discuss global questions, it seems easiest to project the xy-plane stereographically onto the Reimann sphere, the unit sphere in x,y,z-space with its center at (0,0,1). C and Γ then map onto regular curves on the sphere which we shall also call C and Γ, and which meet at the North Pole (0,0,2). Thus the point at ∞ must also be regarded as an equilibrium point of the game G. C is now homeomorphic to a circle. So it makes sense to say that one equilibrium point "separates two others, meaning of course that it lies on the segment of C cut off by them, which does not contain the North

Pole. Also we may speak of the 'neighbors' of a given equilibrium point, by which we shall mean those nearest to it, on either side, along C. Clearly, a given equilibrium may fail to have a neighbor (other than NP) on either or both sides, although in the latter case, it will be the only equilibrium point which the game G possesses.

Next, let us assume that C and Γ actually cross one another every time they meet. Then if Γ crosses C from bottom to top at a point E, Γ must cross from top to bottom at each of E's neighbors. Thus the neighbors of an unstable equilibrium (if any) must be stable. And similarly, the neighbors of a stable equilibrium are unstable. From this fact alone, we can completely describe the qualitative behavior of the solutions of (4.1) for many games.

Let us begin by assuming that $\varphi'(y)$ and $\psi'(x)$ are both positive everywhere, and let us examine the solutions near an unstable equilibrium E, flanked by its two neighbors E^+ and E^-. The situation is indicated in figure 10. Here E has been chosen as the origin of coordinates, and the arrows indicate that the direction field

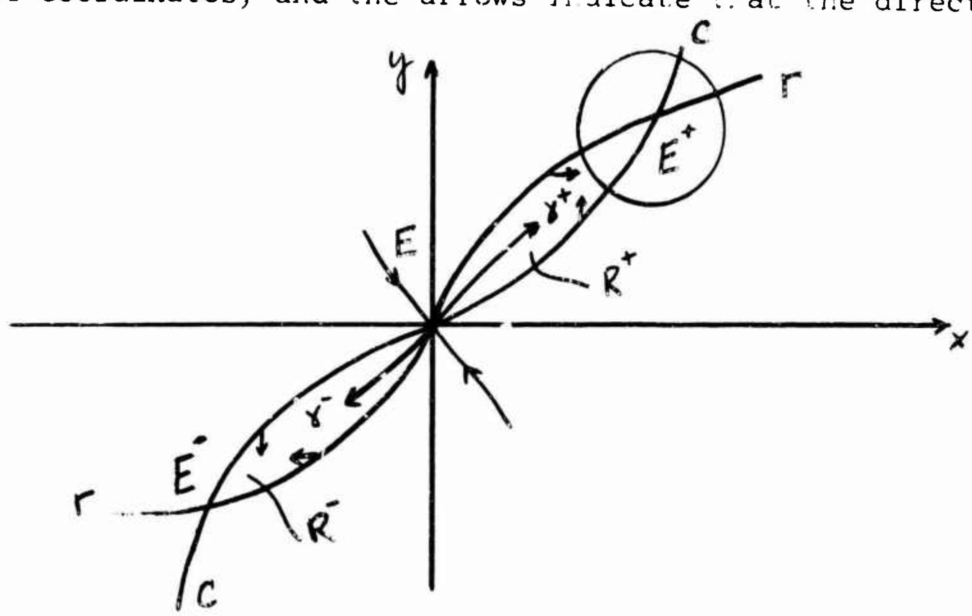

Figure 10.

is horizontal along Γ, pointing toward C, and vertical on C, pointing toward Γ. Thus the regions between C and Γ (labeled R^+ and R^- in the figure) are filled with solutions of positive slope, which are directed away from E. In fact, since E is a saddle point, the solutions near E can only be slightly deformed versions of those in figure 6. In particular, the separatices γ^+ and γ^- must (though they are no longer straight lines) lie in R^+ and R^- respectively, while the incoming separatrices must lie, one each, in the second and fourth quadrants.

Every solution of (4.1) which passes thru a point of R^+ converges to E^+ as t becomes infinite. And in particular, since γ^+ contains points of R^+, γ^+ must converge to E^+. To see this, let K be a circle about E^+ which does not surround E. Let A and B be the points of intersection of K with C and Γ, respectively, and observe that the field vectors at points of $K \cap R^+$ all point into K. And the same is true on any smaller circle as well. Thus no solution that meets the curvilinear triangle E^+AB may ever again leave it. And in this case (see [2], pg. 202), every such solution must tend to the vertex E^+ of the triangle. The assertion concerning γ^+ follows by taking successively larger circles K.

If E has no neighbor on the right, then the image of γ^+ on the Reimann sphere must remain between C and Γ, and converge to NP as t becomes infinite. To see this, let P be a point on γ^+, which lies in R^+, and let KP be the circle about E which passes thru P Balance a unit sphere on E, and project the plane stereographically onto it Then project the sphere, along rays emanating from E, onto the plane z=2 The x and y axes in the original plane will then map onto new x and y axes in z=2, with NP as origin C and Γ will map onto a pair of curves which lie in the first quadrant, and which meet at NP And if we again use A and B to denote the points of intersection of KP with C and Γ, we again have a curvilinear triangle NPAB which γ^+ enters, and in which all solutions tend to a single vertex, namely NP.

Finally, by using the negative x and positive y axes in the roles of C and Γ, and replacing the time parameter t by -t in the system (4 1), we may repeat the above argument to show that the other two separatrices (the ones which end at E) must tend to NP as t approaches $-\infty$

Thus a game with 2n equilibrium points, at each of which C and Γ really cross, must have exactly n saddle points And if C and Γ have everywhere positive slopes, the separatrices which connect the saddle points to NP partition the plane into n+1 strips, of which all but one contain a single stable equilibrium. The situation for n=2 is indicated schematically in figure 11.

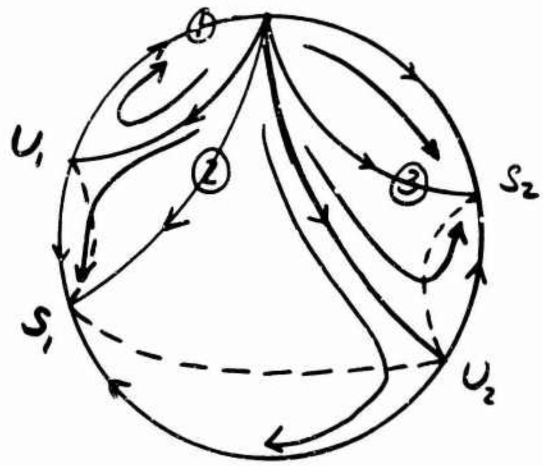

Figure 11.

Here we have taken all of the equilibria to lie on a single great circle in the Reimann sphere, which also passes thru NP , and shown only one of the hemispheres so formed. Region (1) corresponds to a semi-infinite strip, cut off by the separatrices joining NP to the unstable equilibrium U_1 . It contains no stable equilibrium. Region (3) is also semi-infinite, and is cut off by the separatrices from NP to U_2 . It contains the single stable equilibrium S_2 . Region (2) comprises the rest of the plane (sphere), and contains S_1 .

The disparate roles played by the stable and unstable equilibria under conditions of uncertainty are now clear. The stable equilibrium points are the strategy pairs toward which the players will be driven if they follow the procedure of continuous experimentation and strategy modification outlined earlier, while the unstable ones serve to partition the plane into regions from which the various stable equilibria may be reached. Thus if the players begin the game G by playing a strategy pair $(\mathbf{x}_0, y_0)$ in Region (1), and thereafter alter their strategies according to the rules (- 1), they will be led in time to play strategies approximating $x = -$, $y = -$. But if they start in (2) or (3), they will be led eventually to play near S_1 or S_2. And only if they start on one of the separatrices joining NP to U_1 or U_2 will they be led to any other strategy pairs. The situation for larger r is of course similar, but there are more strips and more stable equilibria.

Finally, there are two ways we can make up a game with $2n-1$ equilibria from one with $2n$ without violating our assumption that C and Γ always cross when they meet. We can add an extra stable equilibrium S_0 between U_1 and NP, or we can add an extra unstable one U_{n+1} between S_n and NP.

Of course, we can always add an arbitrary number of equilibria at which C and Γ meet but do not cross, between any pair of neighboring equilibrium points at which they do cross. But it is our feeling that such equilibria are less important than the other kinds, because they would never be observable in practice. So we shall term them "inessential," and confine our attention to games having only essential equilibrium points.

The reason that inessential equilibria can never be observed, of course, is that in reality we can perform only finite experiments. So the functions f and g can only be approximately known, and it can never be determined whether the curves C and Γ actually meet at a point, with out crossing, or just lie very close to one another there. In any case, we shall not discuss inessential equilibria further.

Our results thus far show that the qualitative properties of the solutions of (4.1), for a particular game G, are completely determined* by the number of stable and unstable equilibria G has, provided that (a) the equilibria are all essential, and (b) that C and Γ are the graphs of increasing functions. The latter would be the case, for instance, in any market game of the sort considered in §1, if it were known that whenever firm X increases its price from p_o to p_1, the new optimal price $q_1 = \psi(p_1)$ for Y is higher than the old optimum $q_o = \psi(p_o)$.

There are, of course, many interesting games for which (b) does not hold. And we should like very much to obtain the complete phase-portraits for these games as well. But without (b), or something to replace it, there is no apparent reason that the separatrices emanating from an unstable equilibrium E (see Figure 10) must end at the neighbors E^+ and E^- of E, rather than at any other equilibria. Indeed it is not even clear that they must end at stable equilibria, although we have been unable to produce an example for which they do not. So it is our conjecture that the number of possible global configurations grows rapidly with n, and is probably too large to allow for a simple enumeration like the one we have obtained above under assumption (b). In any event, we shall leave our discussion of global behavior at this point, and turn to other matters.

---

*For we have determined what Lefschetz [ ] calls the "complete phase portrait" of the system.

But first we should like to argue briefly that properties such as I-stability are the "right" ones to study. For there is no reason why X might not adjust $x(t)$ by solving $\dot{x} = 2 f_x$ instead of the system (4.1). And in general, X might choose a "scale function" $r(x,y)$ while Y chooses $\rho(x,y)$, so that the strategy pairs $(x(t), y(t))$ would satisfy

$$(5.1) \qquad \begin{aligned} x &= r(x,y)\, f_x(x,y) \\ \dot{y} &= \rho(x,y)\, g_y(x,y) \end{aligned}$$

But our results are not affected by this! For all we have established about the system (4.1) follows from the assumptions (a) and (b) on the rational curves C and Γ, which are unchanged by the introduction of the scale functions r and ρ. In our view, this observation adds greatly to the "robustness" of our conclusions.

6. Rational Processes

In addition to the processes π_L, π_S, and π_I discussed earlier, there is at least one other process which deserves mention. This we may call the "discrete rational process", because of the essential use it makes of the rational curves C and Γ. To employ it, one player (say X) names a strategy x_o, and allows Y to choose his optimal strategy $y_o = \psi(x_o)$ against it. X then names a second strategy x_1, differing from x_o by no more than ϵ, such that $f(x_1, \psi(x_1)) > f(x_o, \psi(x_o))$. The process terminates when X can no longer find an x_{n+1} such that $f(x_{n+1}, \psi(x_{n+1})) > f(x_n, \psi(x_n))$, and in that case $(x_n, \psi(x_n))$ is obviously an equilibrium point for the game G. We denote the above-described process by the symbol $\pi_R(X;\epsilon)$. A typical process path consists of the line $x = x_o$, plus a polygonal approximation to a portion of Γ, leading from the point where $x = x_o$ meets Γ, to some stable equilibrium point E. We shall sketch a proof of this fact shortly. But before we do, we wish to point out that π_R also has an infinitessimal analogue π_{RI}.

This infinitessimal process consists simply of a straight line beginning somewhere on $x = x_o$, followed by a curve which follows along Γ in the direction of increasing $f(x, \psi(x))$. And it too

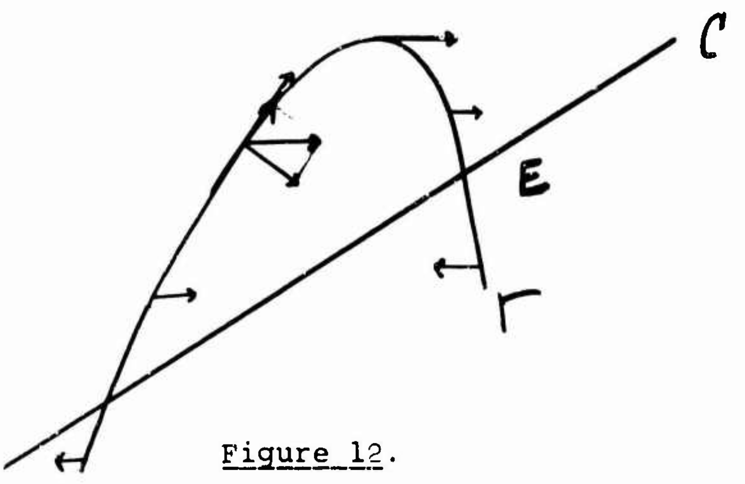

Figure 12.

always converges to a stable equilibrium. To see that this is so, observe that, as shown in Figure 12, there is a well defined direction of increasing $f(x, \psi(x))$ at every point $(x, \psi(x))$ of Γ that is not also a point of C. For if we consider such a point of Γ, and recall that the field vector $(f_x(x,\psi(x)),\ g_y(x,\psi(x)))$ is horizontal there, it is clear that the vector may be resolved into components which are respectively normal and tangential to Γ. And furthermore, the tangential component cannot vanish unless either the whole field vector (f_x,g_y) vanishes, or Γ is vertical at $(x,\psi(x))$. But we saw in §3. that Γ is never vertical, and that the field vector vanished only on $C \cap \Gamma$. Thus the component of (f_x,g_y) which is tangent to Γ may vanish (-i e. change direction) only when $(x,\psi(x))$ crosses C.

Finally, the field vector (f_x,g_y) must point toward the stable equilibrium E in all sufficiently small neighborhoods $N(E)$. And so, therefore, must its tangential component, at points of $\Gamma \cap N(E)$. But if that component can change direction only at points of $C \cap \Gamma$, it must continue to point along Γ toward E, not only in $N(E)$, but also at all points of Γ between the neighbors E^+ and E^- of E. This completes the proof that the process $\pi_{RI}(X)$ must converge to a stable equilibrium point of G, unless the line $x = x_o$ passes thru an unstable equilibrium. In that case, the process consists of a segment of $x = x_o$ only, and converges to the unstable equilibrium which lines thereon.

Similar remarks can of course be made for $\pi_{RI}(Y)$, and the corresponding convergence proofs for the methods $\pi_R(X;\epsilon)$ and $\pi_R(Y;\epsilon)$ follow from the fact their process paths lie within ϵ-neighborhoods of the π_{RI} paths. We point out too that methods quite similar to the methods π_R have been used with some success [4] to compute optimal pursuit and evasion strategies.

7. Constraints

Most of the classical work in game theory has assumed that the strategies x and y were to be chosen from compact convex sets X and Y. In our case, these must each be contained in R, so we may as well take them both to be the unit interval $I = [0,1]$. And we shall refer to the game

$$(G_I) \qquad \max_{x \in I} f(x,y) \qquad\qquad \max_{y \in I} g(x,y)$$

as the "restriction of G to $I \times I$", or simply "the restriction of G".

Clearly, any points of $C \cap \Gamma$ which happen to lie in $I \times I$ are equilibrium points of G_I as well as of G. But G_I may have other equilibria as well. We may call them "induced equilibrium points", as opposed to the "natural equilibria" which lie in $C \cap \Gamma$. The induced equilibria may be of several types.

To begin with, Γ must cross each of the lines $x = 0$ and $x = 1$ exactly once, say at $(0,y_0)$ and at $(1,y_1)$. If $0 \leq y_0 \leq 1$, and $0 \leq y_1 \leq 1$, these points are possible equilibria. To test $(0,y_0)$, observe that C must cross the line $y = y_0$ at a point (x_0,y_0), and that $(0,y_0)$ is an induced equilibrium point of G_I if and only if $x_0 < 0$, or equivalently, if $f_x(0,y_0) < 0$. It is, of course, a natural equilibrium if $x_0 = 0$, and it is not an equilibrium if $x_0 > 0$. Similarly, to see if $(1,y_1)$ is an induced equilibrium point, let (x_1,y_1) be the point where C meets $y = y_1$. Then if $x_1 \geq 1$, $(1,y_1)$ is an equilibrium point of G_I. No other equilibria can lie on the vertical sides of $I \times I$, save possibly at the corners.

There are, similarly, at most two possible equilibrium points on the horizontal sides of IxI, and these may also be tested in a straightforward manner. Either, both, or neither may turn out actually to be equilibria.

If $y_o < 0$, then (0,0) but not (0,1) may be an equilibrium point. To test, look at the point $(x_o, 0)$ at which C crosses $y_o = 0$. If $x_o \leq 0$, then (0,0) is an equilibrium. Otherwise it is not. Similar tests may be performed at the other corners.

Thus G_I may have at most four induced equilibria; no more than one may lie on a single side. In fact, there may never be more than two. For suppose that three (say the top, bottom, and left) sides of IxI contain distinct equilibria. Then neither (0,0) nor (0,1) is an equilibrium. So there must be equilibria at $(x_o, 0)$ and $(x_1, 1)$, where both x_o and x_1 are positive. And if so, the points (x_o, y_o) and (x_1, y_1) at which Γ meets the lines $x = x_o$ and $x = x_1$ must lie below and above IxI, respectively. But Γ must also pass thru the equilibrium point $(0, y^*)$ which lies on the left side of IxI. Hence Γ must either cross the line $x = 0$ twice, or be tangent to it at $(0, y^*)$. And either of these possibilities would contradict the fact that $g_{yy} < 0$. The reader is invited to construct examples having 1 and 2 induced equilibria. How many stable and unstable natural equilibria must such examples posess? And can an induced equilibrium point ever be unstable?

We intend to return to these and related questions at a later date. But for the moment, we are content to suppose that we have revealed something, perhaps quite unexpected, about the set of Nash equilibrium points of a game and the relative ease with which they may be computed. And we hope that others will try the methods we have recommended (especially the rational processes π_R, for which we have great hopes), and find them useful. For we expect that, particularly for games with imperfect information, they are both natural methods to try, and as likely as any to achieve success.

<u>APPENDIX:</u> We present a proof of the lemma in Section 3. The condition (3.5) guarantees that t_0 and t_1 exist (perhaps $\pm\infty$) and that for sufficiently small positive t we have $h(t) < t$ and $h(-t) > -t$. Since $h(t) \neq t$ for all $t \in I^- \cup I^+$, part (i) is proved for $n=1$. A similar argument shows that $h(h(t)) < t$ if $t \in I^+$ and $h(h(t)) > t$ if $t \in I^-$.

We now show $h(t_0) = t_0$ or $h(t_0) = t_1$. By the definition of t_0 we have $h(t_0) \notin I$. If $h(t_0) < t_0$ then, by the continuity of $h(t)-t$, there exists $t \in I^-$ for which $h(t) = t$, a contradiction. So $h(t_0) = t_0$ or $h(t_0) \geq t_1$ and similarly $h(t_1) = t_1$ or $h(t_1) \leq t_0$. Suppose $h(t_0) \geq t_1$. Then there exists $t' \in I^-$ for which $h(t') = t_1$. If $h(t_1) = t_1$ then $h(h(t')) = h(t')$, a contradiction; on the other hand, if $h(t_1) \leq t_0$ then $h(h(t')) \leq t_0 < t'$ so that by the continuity of $h(h(t))-t$ there exists $t'' \in I^-$ for which $h(h(t'')) = t''$, a contradiction. This proves part (ii).

Suppose there exists $t \in I$ for which $h(t) = t_o$. By part (i) with $n=1$ we have $t \in I^+$. If $h(t_o) = t_1$ then $h(h(t)) > t$, a contradiction; if $h(t_o) = t_o$, then $h(h(t)) = h(t)$, a contradiction. A similar argument shows we cannot have $h(t) = t_1$ for $t \in I$. Hence (iii) is proved for $n=1$ and follows for general n by induction.

The case $n=1$ having already been proved, we use induction to prove (i) by assuming it true for $n=1,2,\ldots,m-1$. Suppose $t \in I^+$ and $h^{[m]}(t) = t$. If $h^{[j]}(t) < 0$ for all $j=1,2,\ldots,m-1$, then $h^{[m-1]}(h^{[m-1]}(t)) = h^{[m-2]}(h^{[m]}(t)) = h^{[m-2]}(t) < h^{[m-1]}(t)$, a contradiction; on the other hand, if $h^{[j]}(t) > 0$ for some $j \leq m-1$, then $t = h^{[m]}(t) = h^{[m-j]}(h^{[j]}(t)) < h^{[j]}(t) < t$, a contradiction. A similar argument holds for $t \in I^-$.

To prove (iv) fix t and consider the subsequence $\{h^{[n_j]}(t) : h^{[n_j]}(t) \geq 0\}$. This is a monotone decreasing sequence and hence converges to some $\bar{t} \geq 0$. Suppose $\bar{t} > 0$. Then the subsequence is eventually in any neighborhood of $\bar{t}$. But by the continuity of $h^{[i]}$, if t is in a small neighborhood of $\bar{t}$, $h^{[i]}(t)$ is in a small neighborhood of $h^{[i]}(\bar{t}) < \bar{t}$ so the subsequence cannot remain in any neighborhood of $\bar{t}$. This contradiction proves that the subsequence converges to 0, and a similar argument proves that the complementary subsequence converges to 0 as well.

REFERENCES

1. J. Case, "A Differential Game in Economics," MRCTSR No. 879, Madison, Wisconsin (1...).

S. Lefschetz, "Differential Equations: Geometric Theory," Interscience, New York (1...).

J. B. Rosen, "Existence and Uniqueness of Equilibrium Points for Concave N-Person Games," Econometrica, Vol. 33, No. 3 (196...).

R. D. Turner, "A Heuristic Algorithm for Approximating Min-Max Strategies," Proc. First Intl. Conf. on Th. and Appl. of Diff'l Games; Sept. ...-Oct. 1, 1... at Amherst, Mass.

Unclassified

Security Classification

DOCUMENT CONTROL DATA - R&D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1 ORIGINATING ACTIVITY (Corporate author) Econometric Research Program [illegible]inceton University	2a REPORT SECURITY CLASSIFICATION Unclassified 2b GROUP none

3 REPORT TITLE

ON NASH EQUILIBRIUM POINTS AND GAMES OF IMPERFECT INFORMATION

4 DESCRIPTIVE NOTES (Type of report and inclusive dates)

Summary report: no specific reporting period.

5 AUTHOR(S) (Last name, first name, initial)

James H. Case

6 REPORT DATE	7a TOTAL NO. OF PAGES	7b NO. OF REFS
June 1[illegible]70	[illegible]	[illegible]

8a CONTRACT OR GRANT NO. (N0001[illegible]-[illegible]7 A-01[illegible]1-0007 b PROJECT NO. Task No. 0[illegible]-0[illegible]) c d	9a ORIGINATOR'S REPORT NUMBER(S) Research Memorandum No. 11[illegible] 9b OTHER REPORT NO(S) (Any other numbers that may be assigned this report)

10 AVAILABILITY/LIMITATION NOTICES

Distribution of this document in unliminted.

11 SUPPLEMENTARY NOTES	12 SPONSORING MILITARY ACTIVITY Logistics and Mathematical Branch Office of Naval Research Washington, D.C. [illegible]03[illegible]0

13 ABSTRACT

The importance of Nash equilibrium solutions for certain gaimes of imperfect information is illustrated by means of an example. And motivated thereby, a large number of convergence techniques for the location of Nash equilibria are described and contrasted, for the class of "convex planar games". It is our hope that the availability of this information will both aid and encourage those interested in the Nash equilibria of more complex games.

DD FORM 1 JAN 64 1473

Unclassified

Security Classification

Unclassified
Security Classification

14. KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Games						
Games of imperfect information						
Convergence techniques						

INSTRUCTIONS

1. ORIGINATING ACTIVITY: Enter the name and address of the contractor, subcontractor, grantee, Department of Defense activity or other organization (*corporate author*) issuing the report.

2a. REPORT SECURITY CLASSIFICATION: Enter the overall security classification of the report. Indicate whether "Restricted Data" is included. Marking is to be in accordance with appropriate security regulations.

2b. GROUP: Automatic downgrading is specified in DoD Directive 5200.10 and Armed Forces Industrial Manual. Enter the group number. Also, when applicable, show that optional markings have been used for Group 3 and Group 4 as authorized.

3. REPORT TITLE: Enter the complete report title in all capital letters. Titles in all cases should be unclassified. If a meaningful title cannot be selected without classification, show title classification in all capitals in parenthesis immediately following the title.

4. DESCRIPTIVE NOTES: If appropriate, enter the type of report, e.g., interim, progress, summary, annual, or final. Give the inclusive dates when a specific reporting period is covered.

5. AUTHOR(S): Enter the name(s) of author(s) as shown on or in the report. Enter last name, first name, middle initial. If military, show rank and branch of service. The name of the principal author is an absolute minimum requirement.

6. REPORT DATE: Enter the date of the report as day, month, year, or month, year. If more than one date appears on the report, use date of publication.

7a. TOTAL NUMBER OF PAGES: The total page count should follow normal pagination procedures, i.e., enter the number of pages containing information.

7b. NUMBER OF REFERENCES: Enter the total number of references cited in the report.

8a. CONTRACT OR GRANT NUMBER: If appropriate, enter the applicable number of the contract or grant under which the report was written.

8b, 8c, & 8d. PROJECT NUMBER: Enter the appropriate military department identification, such as project number, subproject number, system numbers, task number, etc.

9a. ORIGINATOR'S REPORT NUMBER(S): Enter the official report number by which the document will be identified and controlled by the originating activity. This number must be unique to this report.

9b. OTHER REPORT NUMBER(S): If the report has been assigned any other report numbers (*either by the originator or by the sponsor*), also enter this number(s).

10. AVAILABILITY LIMITATION NOTICES: Enter any limitations on further dissemination of the report, other than those imposed by security classification, using standard statements such as:

(1) "Qualified requesters may obtain copies of this report from DDC."

(2) "Foreign announcement and dissemination of this report by DDC is not authorized."

(3) "U. S. Government agencies may obtain copies of this report directly from DDC. Other qualified DDC users shall request through ________________."

(4) "U. S. military agencies may obtain copies of this report directly from DDC. Other qualified users shall request through ________________."

(5) "All distribution of this report is controlled. Qualified DDC users shall request through ________________."

If the report has been furnished to the Office of Technical Services, Department of Commerce, for sale to the public, indicate this fact and enter the price, if known.

11. SUPPLEMENTARY NOTES: Use for additional explanatory notes.

12. SPONSORING MILITARY ACTIVITY: Enter the name of the departmental project office or laboratory sponsoring (*paying for*) the research and development. Include address.

13. ABSTRACT: Enter an abstract giving a brief and factual summary of the document indicative of the report, even though it may also appear elsewhere in the body of the technical report. If additional space is required, a continuation sheet shall be attached.

It is highly desirable that the abstract of classified reports be unclassified. Each paragraph of the abstract shall end with an indication of the military security classification of the information in the paragraph, represented as *(TS)*, *(S)*, *(C)*, or *(U)*.

There is no limitation on the length of the abstract. However, the suggested length is from 150 to 225 words.

14. KEY WORDS: Key words are technically meaningful terms or short phrases that characterize a report and may be used as index entries for cataloging the report. Key words must be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location, may be used as key words but will be followed by an indication of technical context. The assignment of links, roles, and weights is optional.

Unclassified
Security Classification